CSCI 357: Algorithmic Game Theory

Lecture 7: GSP & First Price Auctions

Shikha Singh

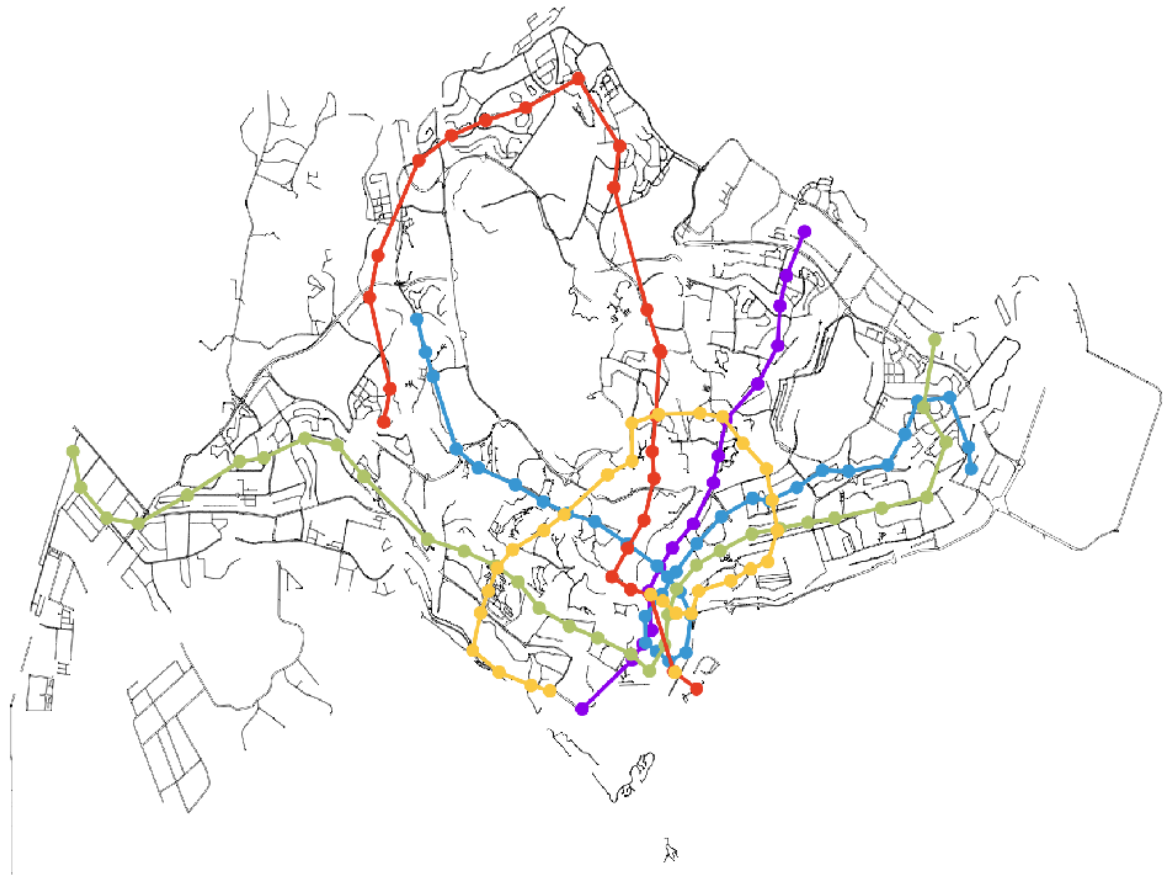

Announcements and Logistics

- **Assignment 3** due by **11 pm tonight**
  - Office hours in this room after class from 4-5.30 pm
  - TA hours may be cancelled tonight: I will hold extra hours either in person/Zoom
- **Assignment 2:** Feedback returned, Q3 assumption in question regraded
  - Collecting student solutions: will be posted on GLOW
- **Assignment 4** will also be a **partner assignment** (simulation based)
  - Everyone must fill out google form by **tomorrow 2 pm** to indicate your partner status and Github ID: https://tinyurl.com/357partner
  - Goal: simulate simple auctions, do empirical analysis of how revenue/utility of agents

Questions?

Last Time

- Started analyzing the **generalized second price auction**
- Defined the Nash equilibrium conditions
  - Too many Nash equilibria: some that are not socially efficient
- Defined a stronger condition: envy-free Nash equilibrium
  - Intuition: no one wants another slot at its current price
- Still many envy-free Nash
  - Started to reason about which envy-free bidders are likely to reach
- Defined balanced bidding strategies:
  - Can also think of them as "locally" envy free

Recall Envy-Free Nash

- The challenge in analyzing GSP is that there can be multiple equilibria
  - How do bidders select---depends on which equilibria is more plausible and reached by a straightforward bidding strategy
- **Envy-free outcome.** We say that a bid profile $\mathbf{b} = (b_1, \ldots, b_n)$ where $b_1 \geq b_2 \geq \ldots \geq b_n$ is **envy-free** if for every bidder i

$$\alpha_i(v_i - b_{i+1}) \geq \alpha_j(v_i - b_{j+1})$$

- **Interpretation:** (current price-per-click of slot j is $p_j = b_{j+1}$)
  - *each bidder i is as happy getting its current slot at its current price as it would be getting any other slot at that slot's current price*

Recall Balanced Bidding

- We say a bid profile $\mathbf{b} = (b_1, b_2, \ldots, b_n)$ satisfies the balanced bidding requirement if
  - The following holds for bidder i for $2 \le i \le m$

$$\underbrace{\alpha_i(v_i - b_{i+1})}_{\text{utility current position}} = \underbrace{\alpha_{i-1}(v_i - b_i)}_{\text{utility in case of retaliation}}$$

  - Any unassigned bidder bids their true value
- Notice that for value ordered bids, the balanced bidding requirement defines **a unique bid profile** (up to the indifference of the top bidder)

Credit: Textbook by Parkes and Seuken

Balanced Bidding Strategies

- b_i must be the highest bid such that bidder i is indifferent between remaining in slot i and having bidder $i-1$ retaliate

$\alpha_1 = 0.2$ $v_1 = 10$ $b_1 = 10$

$\alpha_2 = 0.18$ $v_2 = 4$ $b_2 = ?$

$\alpha_3 = 0.1$ $v_3 = 2$ $b_3 = ?$

$v_3 = 1$ $b_4 = 1$

Balanced Bidding Strategies

- Bid b_3 must be the highest bid such that bidder 3 is indifferent between remaining in slot 2 and having bidder 1 retaliate
- $0.1(2 - 1) = 0.18(2 - b_3) \implies b_3 = ?$

$\alpha_1 = 0.2$	$v_1 = 10$	$b_1 = 10$
$\alpha_2 = 0.18$	$v_2 = 4$	$b_2 = ?$
$\alpha_3 = 0.1$	$v_3 = 2$	$b_3 = 13/9$
	$v_3 = 1$	$b_4 = 1$

Balanced Bidding Strategies

- Bid b_2 must be the highest bid such that bidder 2 is indifferent between remaining in slot 2 and having bidder 1 retaliate:

$$0.18(4 - 13/9) = 0.2(4 - b_2) \implies b_2 = 17/10$$

$\alpha_1 = 0.2$ $v_1 = 10$ $b_1 = 10$ $p_1 = 17/10$

$\alpha_2 = 0.18$ $v_2 = 4$ $b_2 = 17/10$ $p_2 = 13/9$

$\alpha_3 = 0.1$ $v_3 = 2$ $b_3 = 13/9$ $p_3 = 1$

$v_3 = 1$ $b_4 = 1$

Compare to VCG

- **Exercise:** Compute the VCG payments for this example

These are exactly the VCG payments!!!

$\alpha_1 = 0.2$ $v_1 = 10$ $b_1 = 10$ $p_1 = 17/10$

$\alpha_2 = 0.18$ $v_2 = 4$ $b_2 = 17/10$ $p_2 = 13/9$

$\alpha_3 = 0.1$ $v_3 = 2$ $b_3 = 13/9$ $p_3 = 1$

$v_3 = 1$ $b_4 = 1$

Bigger Picture

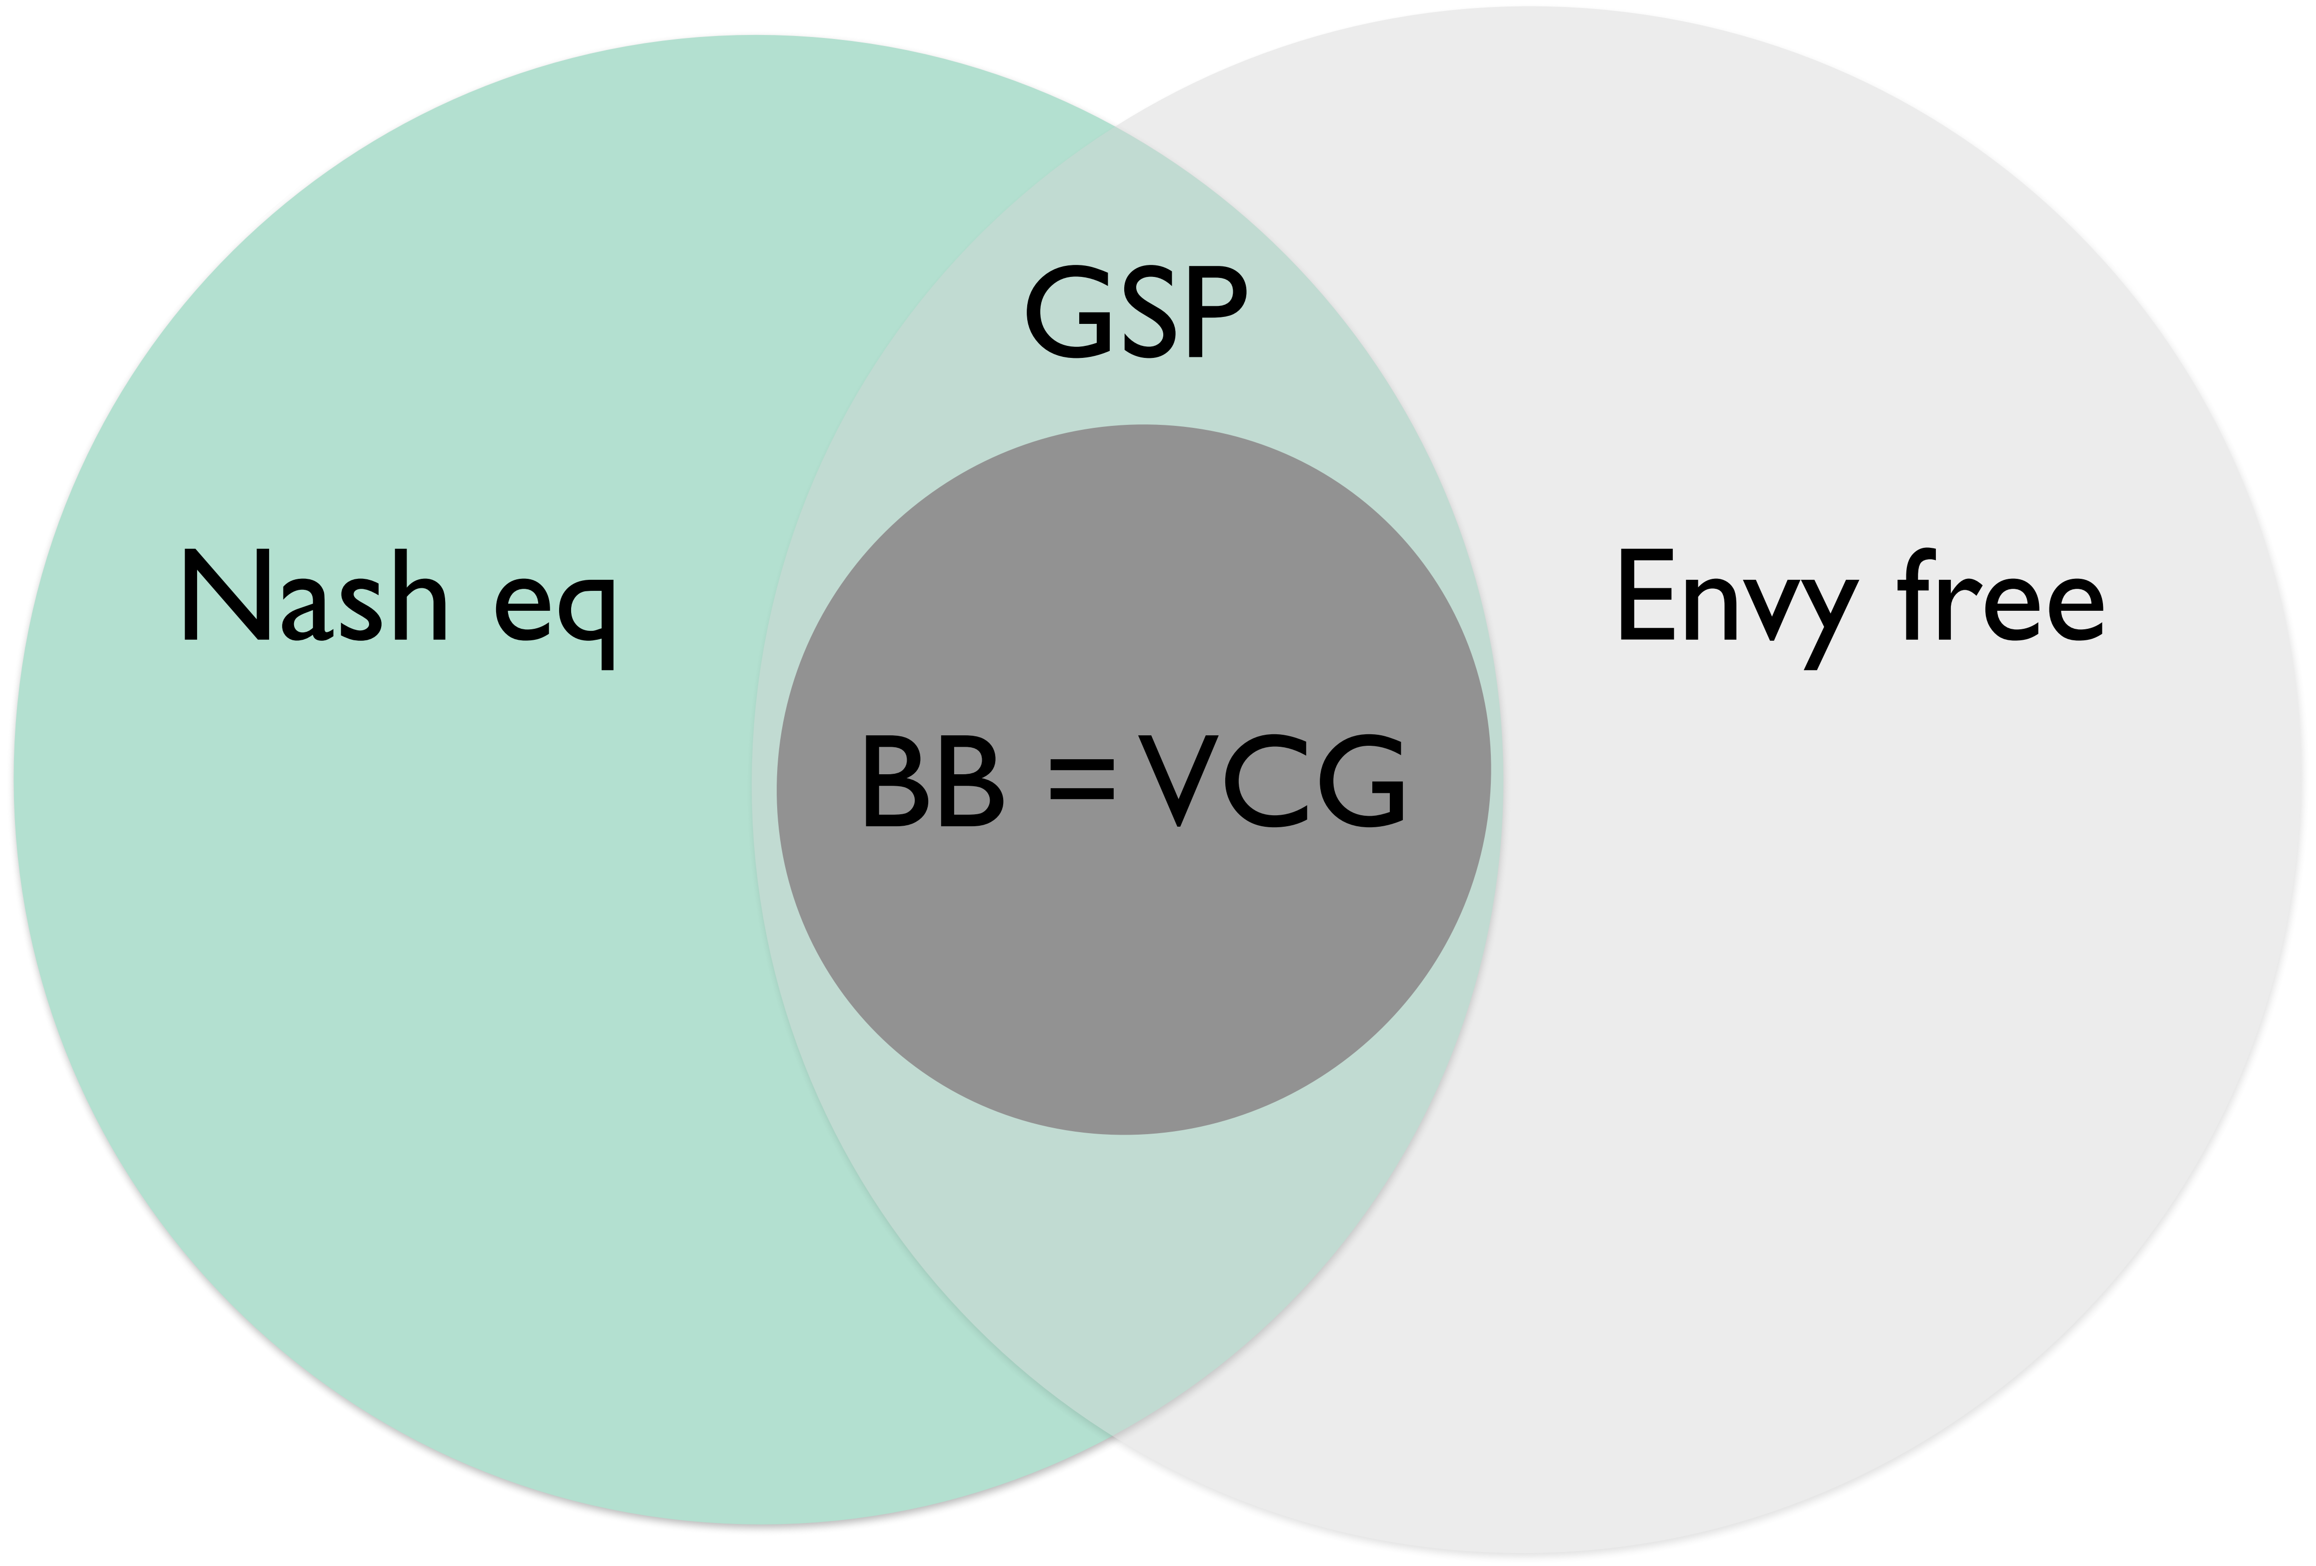

Figure adapted from Textbook by Parkes and Seuken

Balanced Bidding $\Longrightarrow$ Envy Free

Balanced Bidding

- **Lemma.** There exists an envy-free Nash equilibrium of the GSP auction where all bidder's bids are value ordered and satisfy balanced bidding condition.
- **Proof.** Consider a bid profile $\mathbf{b}$ that satisfies balanced bidding with the top bidder bidding their true value $b_1 = v_1$
- Utility of bidder 1 for truthful bidding is $u_1 = \alpha_1(v_1 - b_2)$
- Want to show $u_1 \geq \alpha_2(v_1 - b_3)$

Credit: Textbook by Parkes and Seuken

Proof from the Board

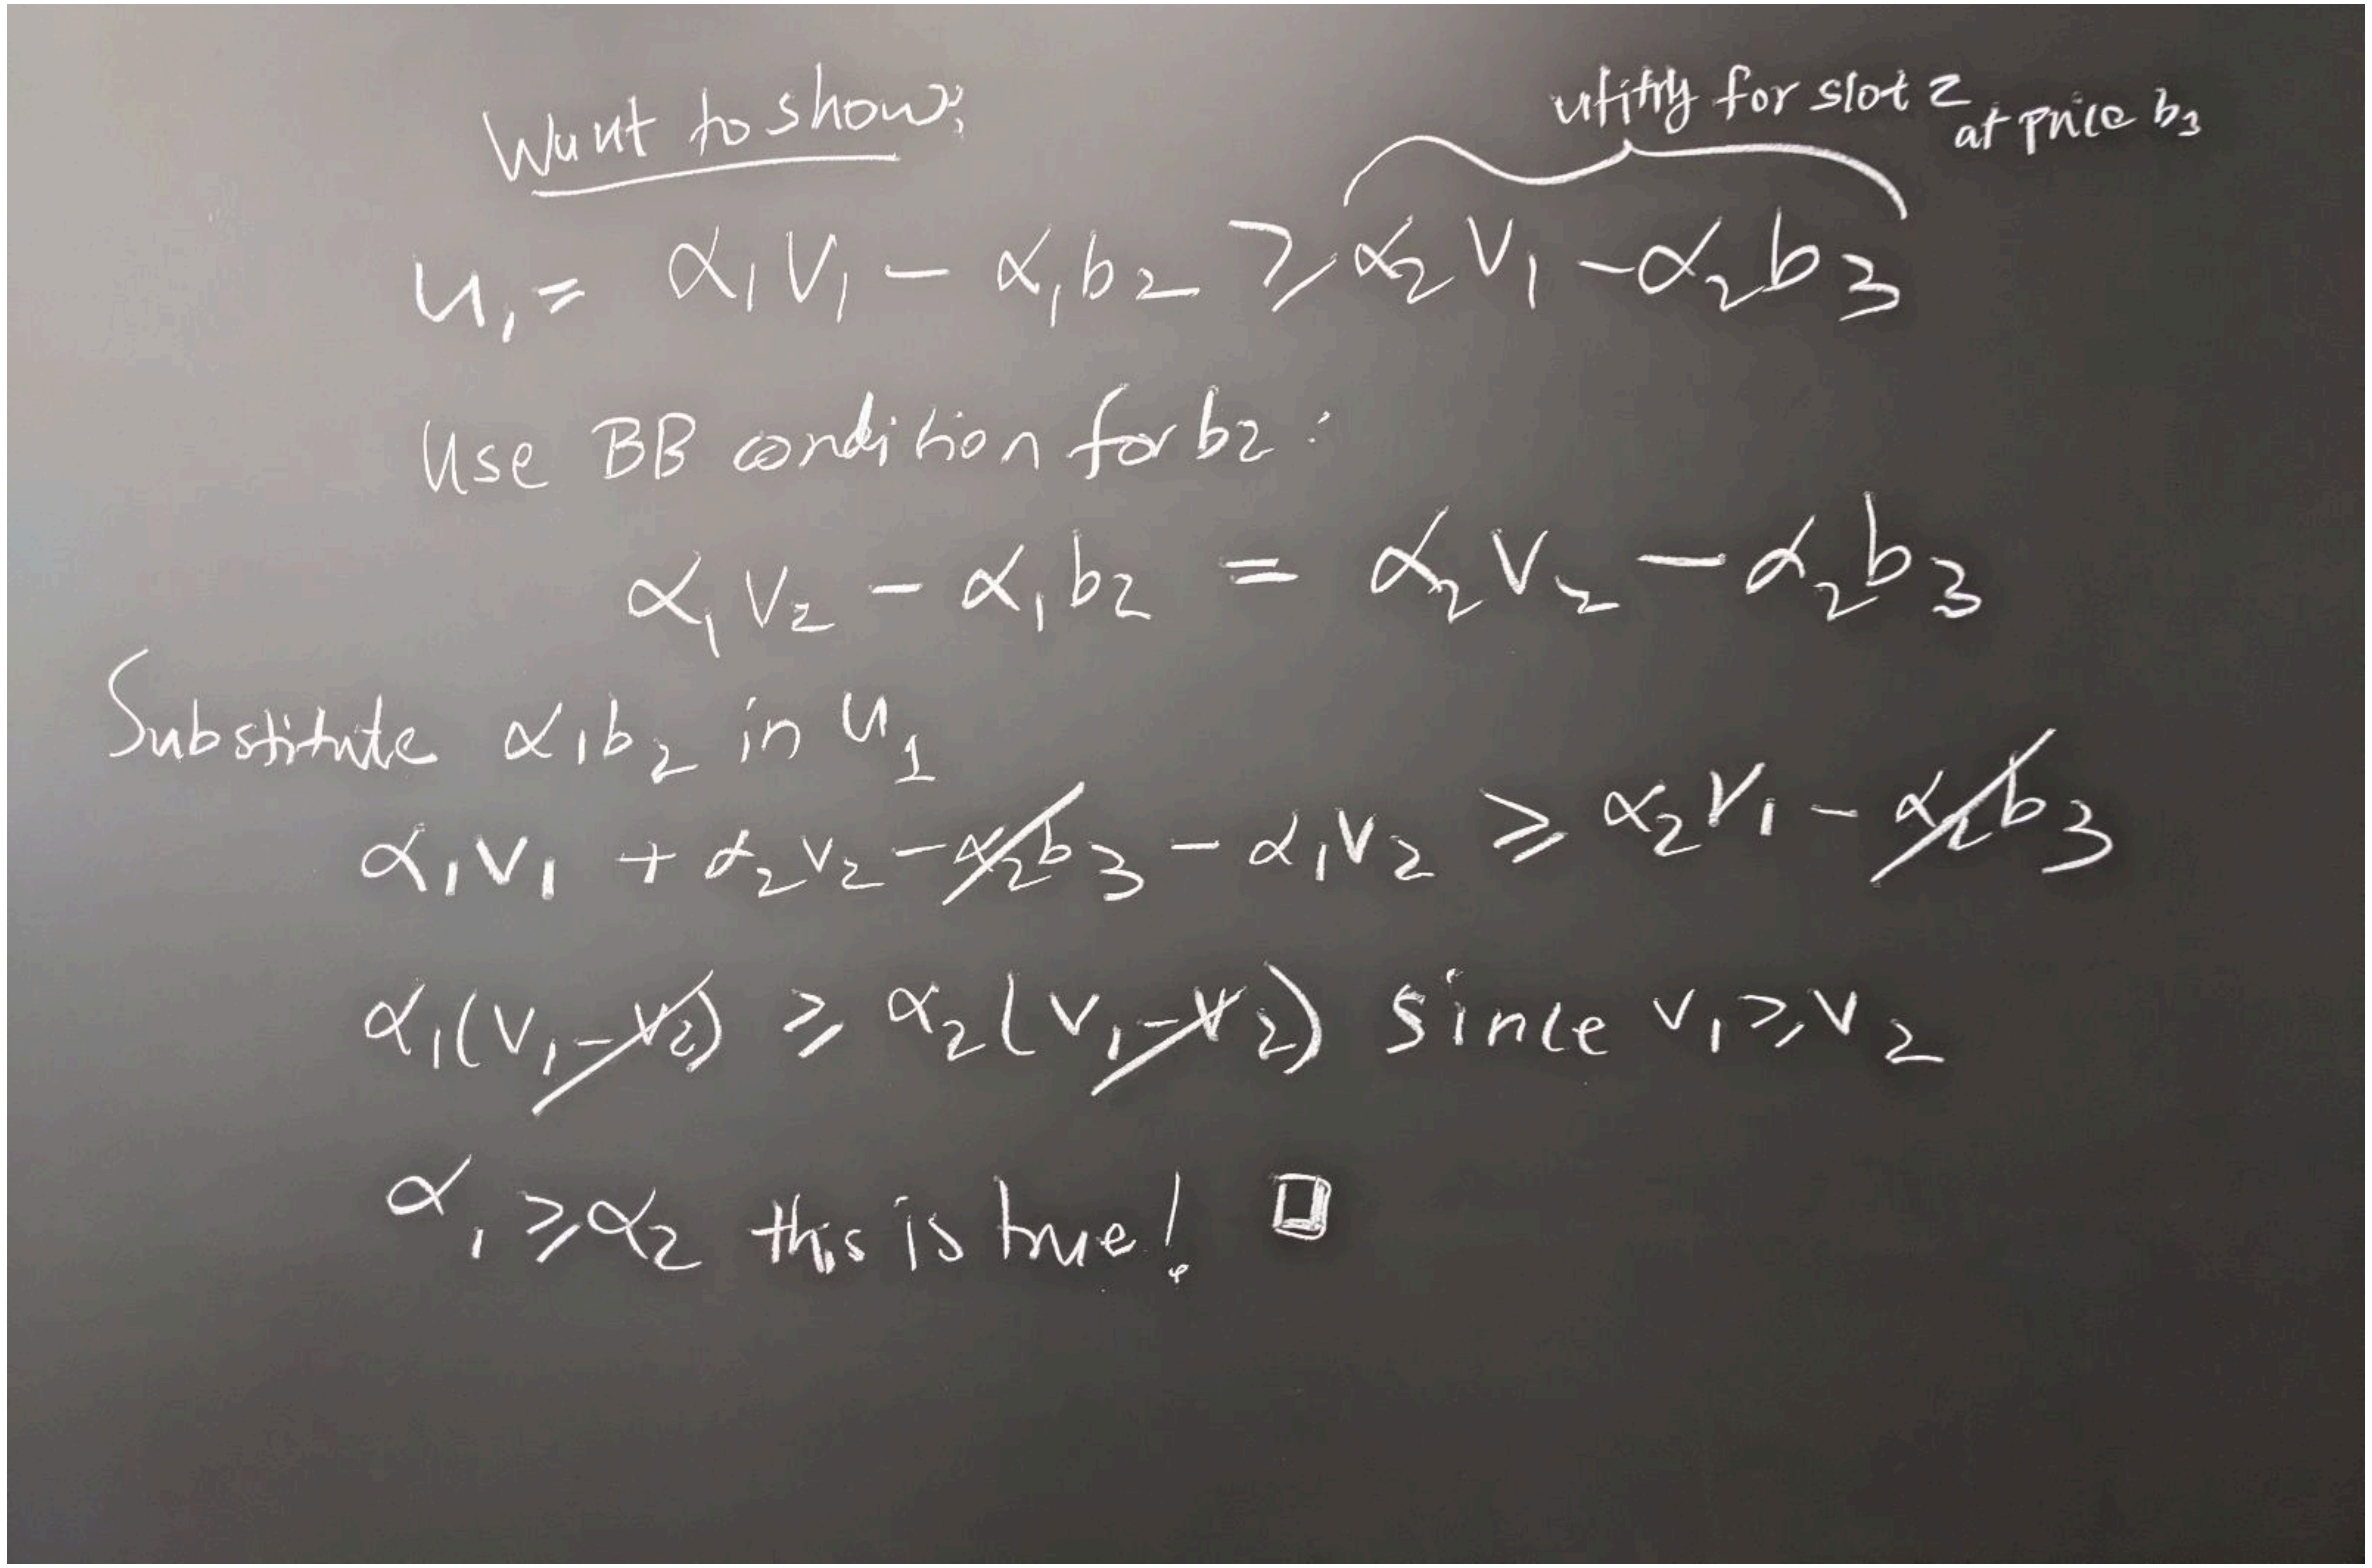

Balanced Bidding

- **Lemma.** There exists an envy-free Nash equilibrium of the GSP auction where all bidder's bids are value ordered and satisfy balanced bidding condition.
- **Proof.** Consider a bid profile $\mathbf{b}$ that satisfies balanced bidding with the top bidder bidding their true value $b_1 = v_1$
- Continuing similarly, we can show that bidder 1 has no utility to deviate to get any lower slot, that is,
  - $\alpha_1(v_1 - b_2) \geq \alpha_2(v_1 - b_3) \geq \ldots \geq \alpha_k(v_1 - b_{k+1})$
- Similarly, we can argue that any bidders $2,\ldots,k$ do not have incentive to deviate to get a lower slot
- To finish the proof, we need to argue the same for **upward deviations**

Credit: Textbook by Parkes and Seuken

Balanced Bidding

- **Lemma.** There exists an envy-free Nash equilibrium of the GSP auction where all bidder's bids are value ordered and satisfy balanced bidding condition.
- **Proof.** Consider a bid profile $\mathbf{b}$ that satisfies balanced bidding with the top bidder bidding their true value $b_1 = v_1$
- Consider bidder 3 in position 3
  - Current utility is $\alpha_3(v_3 - b_4)$
  - In balanced bidding we have $\alpha_3(v_3 - b_4) = \alpha_2(v_3 - b_3)$
    - No incentive to deviate to slot 2
  - Similarly, we can use BB condition: $\alpha_1(v_2 - b_2) = \alpha_2(v_2 - b_3)$
  - $\alpha_1(v_3 - b_2) \ \leq \alpha_3(v_3 - b_4)$ (no incentive to target slot 1)
- We can generalize this to any bidder making upward deviations ∎

Credit: Textbook by Parkes and Seuken

(GSP, Balanced Bidding) $\equiv$
(VCG with Truthful)

Balanced Bidding and VCG

- **Theorem.** The outcome of the GSP auction in an envy-free Nash equilibrium bid profile $\mathbf{b}$ that satisfies balanced bidding is equal to the truthful outcome of the VCG auction.
- **Proof.** As the bids are value ordered in an envy-free Nash equilbrium, the allocation in GSP is the same as VCG (and thus surplus maximizing)
- To show that payments are equivalent, we use induction
- We can write the VCG payment (under truthful bids) recursively as:
  - $p_i[\text{VCG}] = (\alpha_i - \alpha_{i+1})v_{i+1} + p_{i+1}(v)$ for $1 \leq i \leq k$
- Let $p_i[\text{GSP}] = \alpha_i b_{i+1}$ be the GSP payment of bidders $1 \leq i \leq k$
- For $i > k$, both mechanisms charge zero.
- We can show $p_i[\text{VCG}] = p_i[\text{GSP}]$ using induction.
- **Base case** $i = k$. $p_k[\text{VCG}] = p_k[\text{GSP}] = \alpha_k b_{k+1}$

Credit: Textbook by Parkes and Seuven

Balanced Bidding and VCG

- **Theorem.** The outcome of the GSP auction in an envy-free Nash equilibrium bid profile $\mathbf{b}$ that satisfies balanced bidding is equal to the truthful outcome of the VCG auction.
- **Proof.**
- We can show $p_i[\text{VCG}] = p_i[\text{GSP}]$ using induction.
- Base case $i = k$. $p_k[\text{VCG}] = p_k[\text{GSP}] = \alpha_k b_{k+1}$
- For a slot $i < k$, we have: $p_i[\text{GSP}] = \alpha_i b_{i+1}$
  - Applying balanced bidding on bidder $i + 1$, we get $\alpha_{i+1}(v_{i+1} - b_{i+2}) = \alpha_i(v_{i+1} - b_{i+1})$
- That is, $p_i[\text{GSP}] = \alpha_i b_{i+1} = (\alpha_i - \alpha_{i+1})v_{i+1} + \alpha_{i+1} b_{i+2}$
  - $= (\alpha_i - \alpha_{i+1})v_{i+1} + p_{i+1}[\text{GSP}]$
  - $= (\alpha_i - \alpha_{i+1})v_{i+1} + p_{i+1}[\text{VCG}] = p_i[\text{VCG}]$ ∎

Credit: Textbook by Parkes and Seuven

Takeaways

- What are some interesting takeaways from this discussion?

2020 Nobel Prize for Auction Theory

© Nobel Media. Ill. Niklas Elmehed.

Paul R. Milgrom

Prize share: 1/2

© Nobel Media. Ill. Niklas Elmehed.

Robert B. Wilson

Prize share: 1/2

The Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel 2020 was awarded jointly to Paul R. Milgrom and Robert B. Wilson "for improvements to auction theory and inventions of new auction formats."

- Nobel Prize Press Release

https://www.nobelprize.org/prizes/economic-sciences/2020/press-release/

> Their theoretical discoveries have improved auctions in practice

\- Nobel Prize Press Release

https://www.nobelprize.org/prizes/economic-sciences/2020/press-release/

Ad Auctions Landscape

GSP Auction Research

Internet Advertising and the Generalized Second Price Auction: Selling Billions of Dollars Worth of Keywords

Benjamin Edelman
Harvard University
bedelman@fas.harvard.edu

Michael Ostrovsky
Stanford University
ostrovsky@gsb.stanford.edu

Michael Schwarz*
UC Berkeley and NBER
mschwarz@berkeley.edu

October 3, 2005

Abstract

We investigate the "generalized second price" auction (GSP), a new mechanism which is used by search engines to sell online advertising that most Internet users encounter daily. GSP is tailored to its unique environment, and neither the mechanism nor the environment have previously been studied in the mechanism design literature. Although GSP looks similar to the Vickrey-Clarke-Groves (VCG) mechanism, its properties are very different. In particular, unlike the VCG mechanism, GSP generally does not have an equilibrium in dominant strategies, and truth-telling is not an equilibrium of GSP. To analyze the properties of GSP in a dynamic environment, we describe the generalized English auction that corresponds to the GSP and show that it has a unique equilibrium. This is an ex post equilibrium that results in the same payoffs to all players as the dominant strategy equilibrium of VCG.

Design Tradeoffs in Practice

- Which auction format (GSP or VCG) is used in various ad market auctions keeps evolving
- It may seem like VCG should always be the first choice: it is in fact an awesome auction for sponsored search
- However, Google and Bing continue to use GSP. Many reasons for this:
  - Inertia: well, it was implemented first (perhaps by accident)
  - Easy to explain pricing rule
  - Short-term revenue loss for switching: for a fixed bid profile, GSP gives more revenue than VCG so switching would lead to short-term revenue loss
  - Re-engineering cost: search engines run campaigns on behalf of marketers and built considerable infrastructure around GSP

Credit: Textbook by Parkes and Seuven

Design Tradeoffs in Practice

- VCG is used for contextual non real-time advertising, e.g. by Twitter and Facebook
- There has also been recent switch to VCG for others, e.g.
  - The Yandex search engine switched from GSP to GSP in 2015
  - In 2012, Google switched from GSP to VCG for its ad network AdSense
- Reasons to prefer VCG over GSP
  - Truthful behavior: no need for bidders to strategize
  - Easier for sellers to estimate revenue
  - Enables faster experimentation: seeing how reserve prices effect revenue, etc.
  - Flexibility: VCG auction is highly configurable to different preferences and contexts

Credit: Textbook by Parkes and Seuven

Design Tradeoffs in Practice

"With respect to flexibility, in 2002, the important decisions were how to rank ads and how to rank ads and how to price ads and the GSP handled these decisions well. By 2012, there were other treatments that could be applied to ads. One particularly useful ad treatment is known as 'dynamic resizing.' It turns out that if you have one highly relevant and three so-so ads, you get more total clicks by enlarging the size of the highly relevant ad and showing it alone."

-- Varian and Harris

VCG in Theory and Practice

Hal R. Varian

Christopher Harris

Google, Inc.

May 2013

Recent Switch to First Price Auctions

- Ad exchanges moved from second-price sealed bid to first-price sealed-bid, with Google switching during **2019**
- **Transparency**. Some businesses are both sellers and buyers
- Composability

	Non real-time	**Real-time (programmatic)**
Sponsored search	◦ Google and Bing **GSP** ◦ Yandex **VCG**	n/a
Contextual	Own inventory ◦ Twitter and Facebook **VCG** 3rd party inventory (ad networks) ◦ Google AdSense, FB Audience Network **VCG** ◦ Microsoft Audience Ads **GSP**	3rd party inventory (ad exchanges) ◦ AppNexus, Twitter, MoPub, and Google DoubleClick **First price** (was second price)

Credit: Textbook by Parkes and Seuven

Why Do Competitive Markets Converge to First-Price Auctions?

Renato Paes Leme
renatoppl@google.com
Google Research
New York, NY

Balasubramanian Sivan
balusivan@google.com
Google Research
New York, NY

Yifeng Teng*
yifengt@cs.wisc.edu
University of Wisconsin-Madison
Madison, WI

ABSTRACT

We consider a setting in which bidders participate in multiple auctions run by different sellers, and optimize their bids for the *aggregate* auction. We analyze this setting by formulating a game between sellers, where a seller's strategy is to pick an auction to run. Our analysis aims to shed light on the recent change in the Display Ads market landscape: here, ad exchanges (sellers) were mostly running second-price auctions earlier and over time they switched to variants of the first-price auction, culminating in Google's Ad Exchange moving to a first-price auction in 2019. Our model and results offer an explanation for why the first-price auction occurs as a natural equilibrium in such competitive markets.

"Moving to a first-price auction puts Google at parity with other exchanges and SSPs in the market, and will contribute to a much fairer transactional process across demand sources." : Scott Mulqueen

Lots of Avenues for Delving Deeper

- Sponsored search markets have been a topic of extensive research in AGT
- If interested, lots of avenues to delve deeper (as part of Final Project)
- Many papers to start "skimming" to find interesting algorithms you want to implement or want to learn the theory behind
  - Greedy Bidding Strategies for Keyword Auctions, Cary et al https://homes.cs.washington.edu/~karlin/papers/ecc.pdf
  - On Revenue in the Generalized Second Price Auction Lucier et al. https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.298.4176&rep=rep1&type=pdf
  - Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords https://www.benedelman.org/publications/gsp-060801.pdf
- Roughgarden's course on **Foundations of Sponsored Search**: http://timroughgarden.org/f07/f07.html (a bit outdated but has useful resources)

Homework 4 Preview

Balanced Bidding Agent

- Implement a "greedy" strategy to design a balanced bidding agent
- Repeated keyword auctions over rounds (model a full day)
- Agents start with an initial then update their bidding
    - Assume bids in next round stay the same as last
    - Target a slot that gives best utility
    - Bid a value that follows balanced bidding conditions
- Does such a "iterative updating strategy" converge?
- Empirical evidence suggests that it does converge in an equilibrium

Greedy Bidding Strategies for Keyword Auctions

Matthew Cary
University of Washington
Aparna Das
Brown University
Ben Edelman
Harvard University
Ioannis Giotis
University of Washington
Kurtis Heimerl
University of Washington
Anna R. Karlin
University of Washington
Claire Mathieu
Brown University
Michael Schwarz
Yahoo! Research

Credit: Textbook by Parkes and Seuven

Ad Auction Simulator

- The auction proceeds in rounds $t \leftarrow 0,1,\ldots,47$
    - Each round models 30 minutes so 48 rounds model a full day
- Value-per-click distribution of the bidding agents are drawn uniformly from 25 cents to 175 cents
- Click-through rates:
    - Top slot $\alpha_1^t = 30\cos(\pi t/24) + 50$
    - Other slots $\alpha_i^t = 0.75^{i-1}\alpha_1^t$

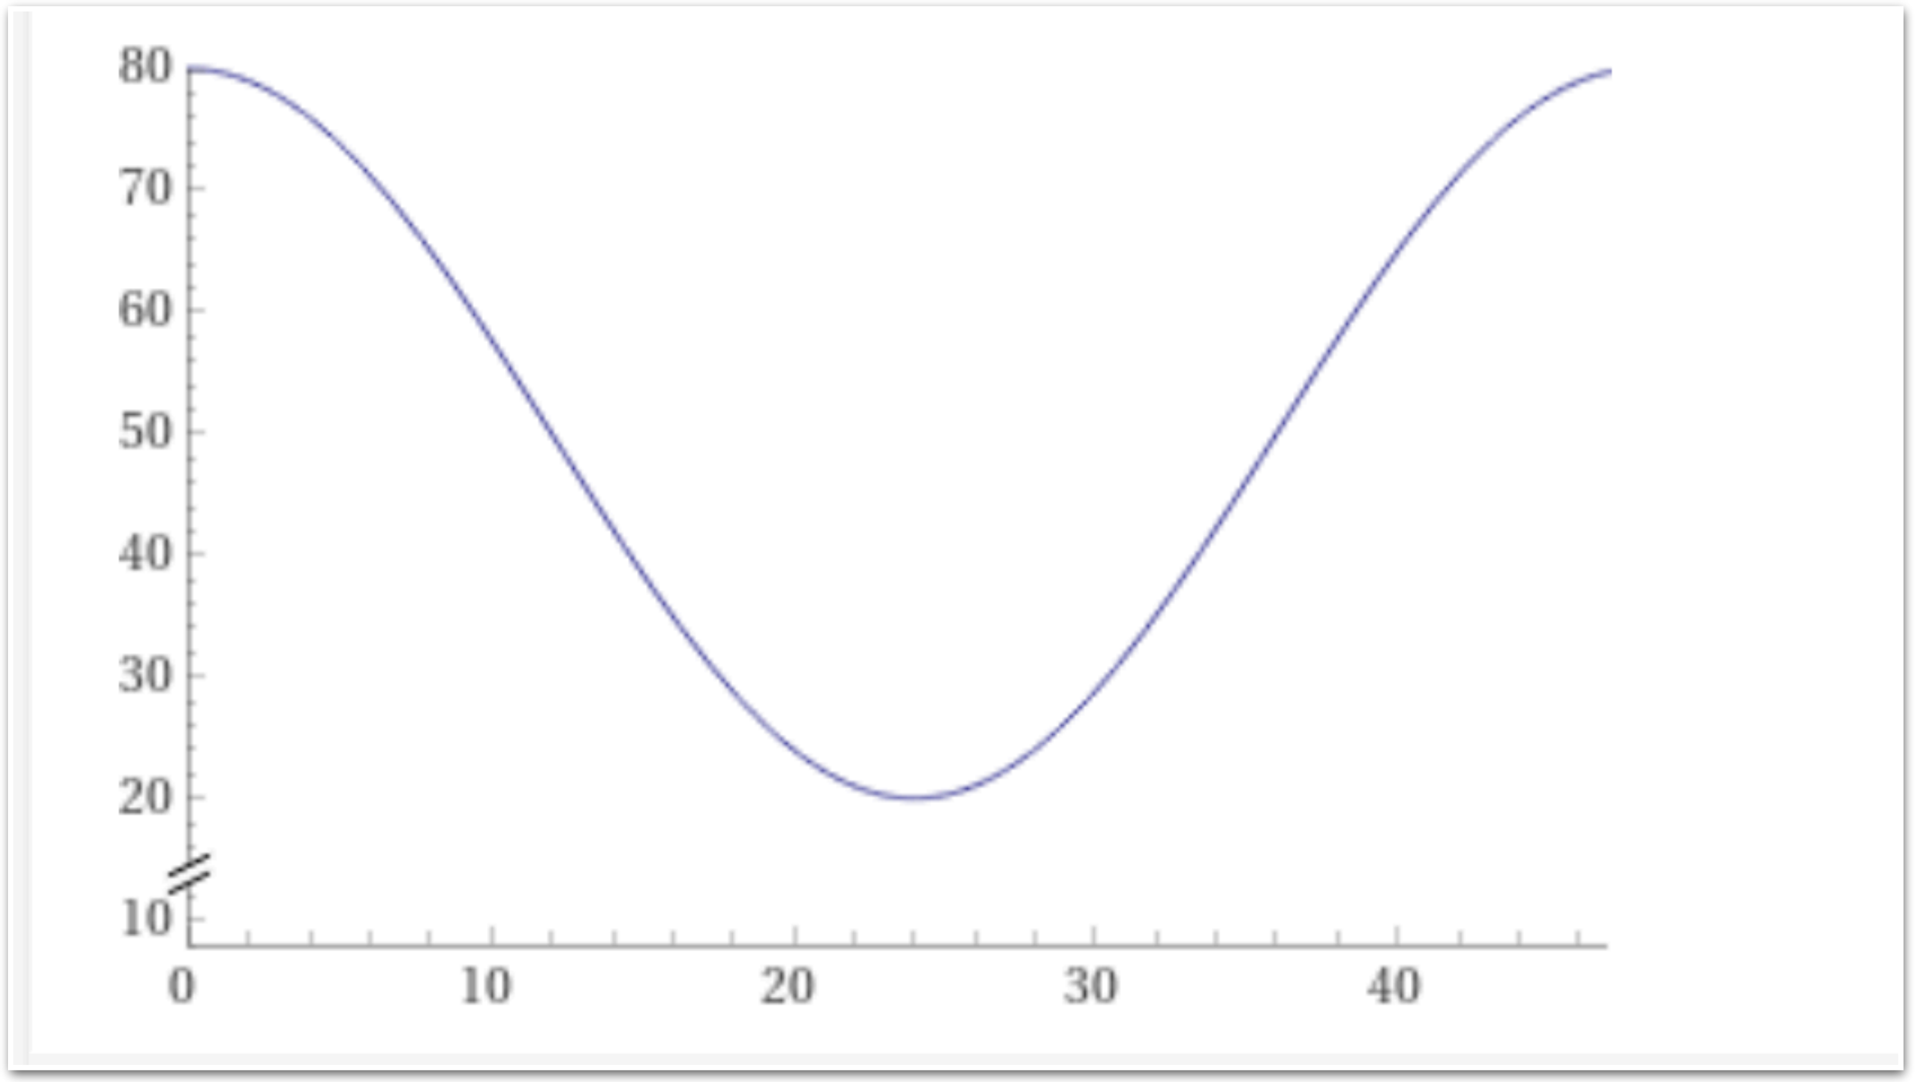

Python3, GitHub and Submission

- Make sure you have Python 3.7 or above installed
  - Guide from 134: https://williams-cs.github.io/cs134-s20-www/shikha-lectures/134-Lecture15.pdf
- Submit code through Git and written part of the homework through Gradescope (final grades would appear on Gradescope)
- Presenting simulation results in write up
  - Use tables and label it with the test parameters used
  - Interpret your results, and explain why you think you are getting the numbers that you are
- **Open ended.** Most of the assignment is open ended evaluations
- **Debugging.** You have to spend time reading the stats and making sure things are working as they should

First Price Auctions

Why are they tricky?

- When bidders have private values are games of **incomplete information**
- In a complete-information game, the following is common knowledge:
  - number of players and actions available to each player
  - the payoff associated with each action profile
- Why was this not a problem when we analyze second-price auction?
  - bidders have a dominant strategy (makes it irrelevant whether or not they know about the values of others to reason about their own)
- Do bidders have a dominant strategy in first-price auction?
- Without a DSE, what is the next best solution concept?
  - Nash equilibrium (but this only works for complete-information games)
- For incomplete-information games: need **Bayesian Nash equilibrium**

Valuations (Sorted)

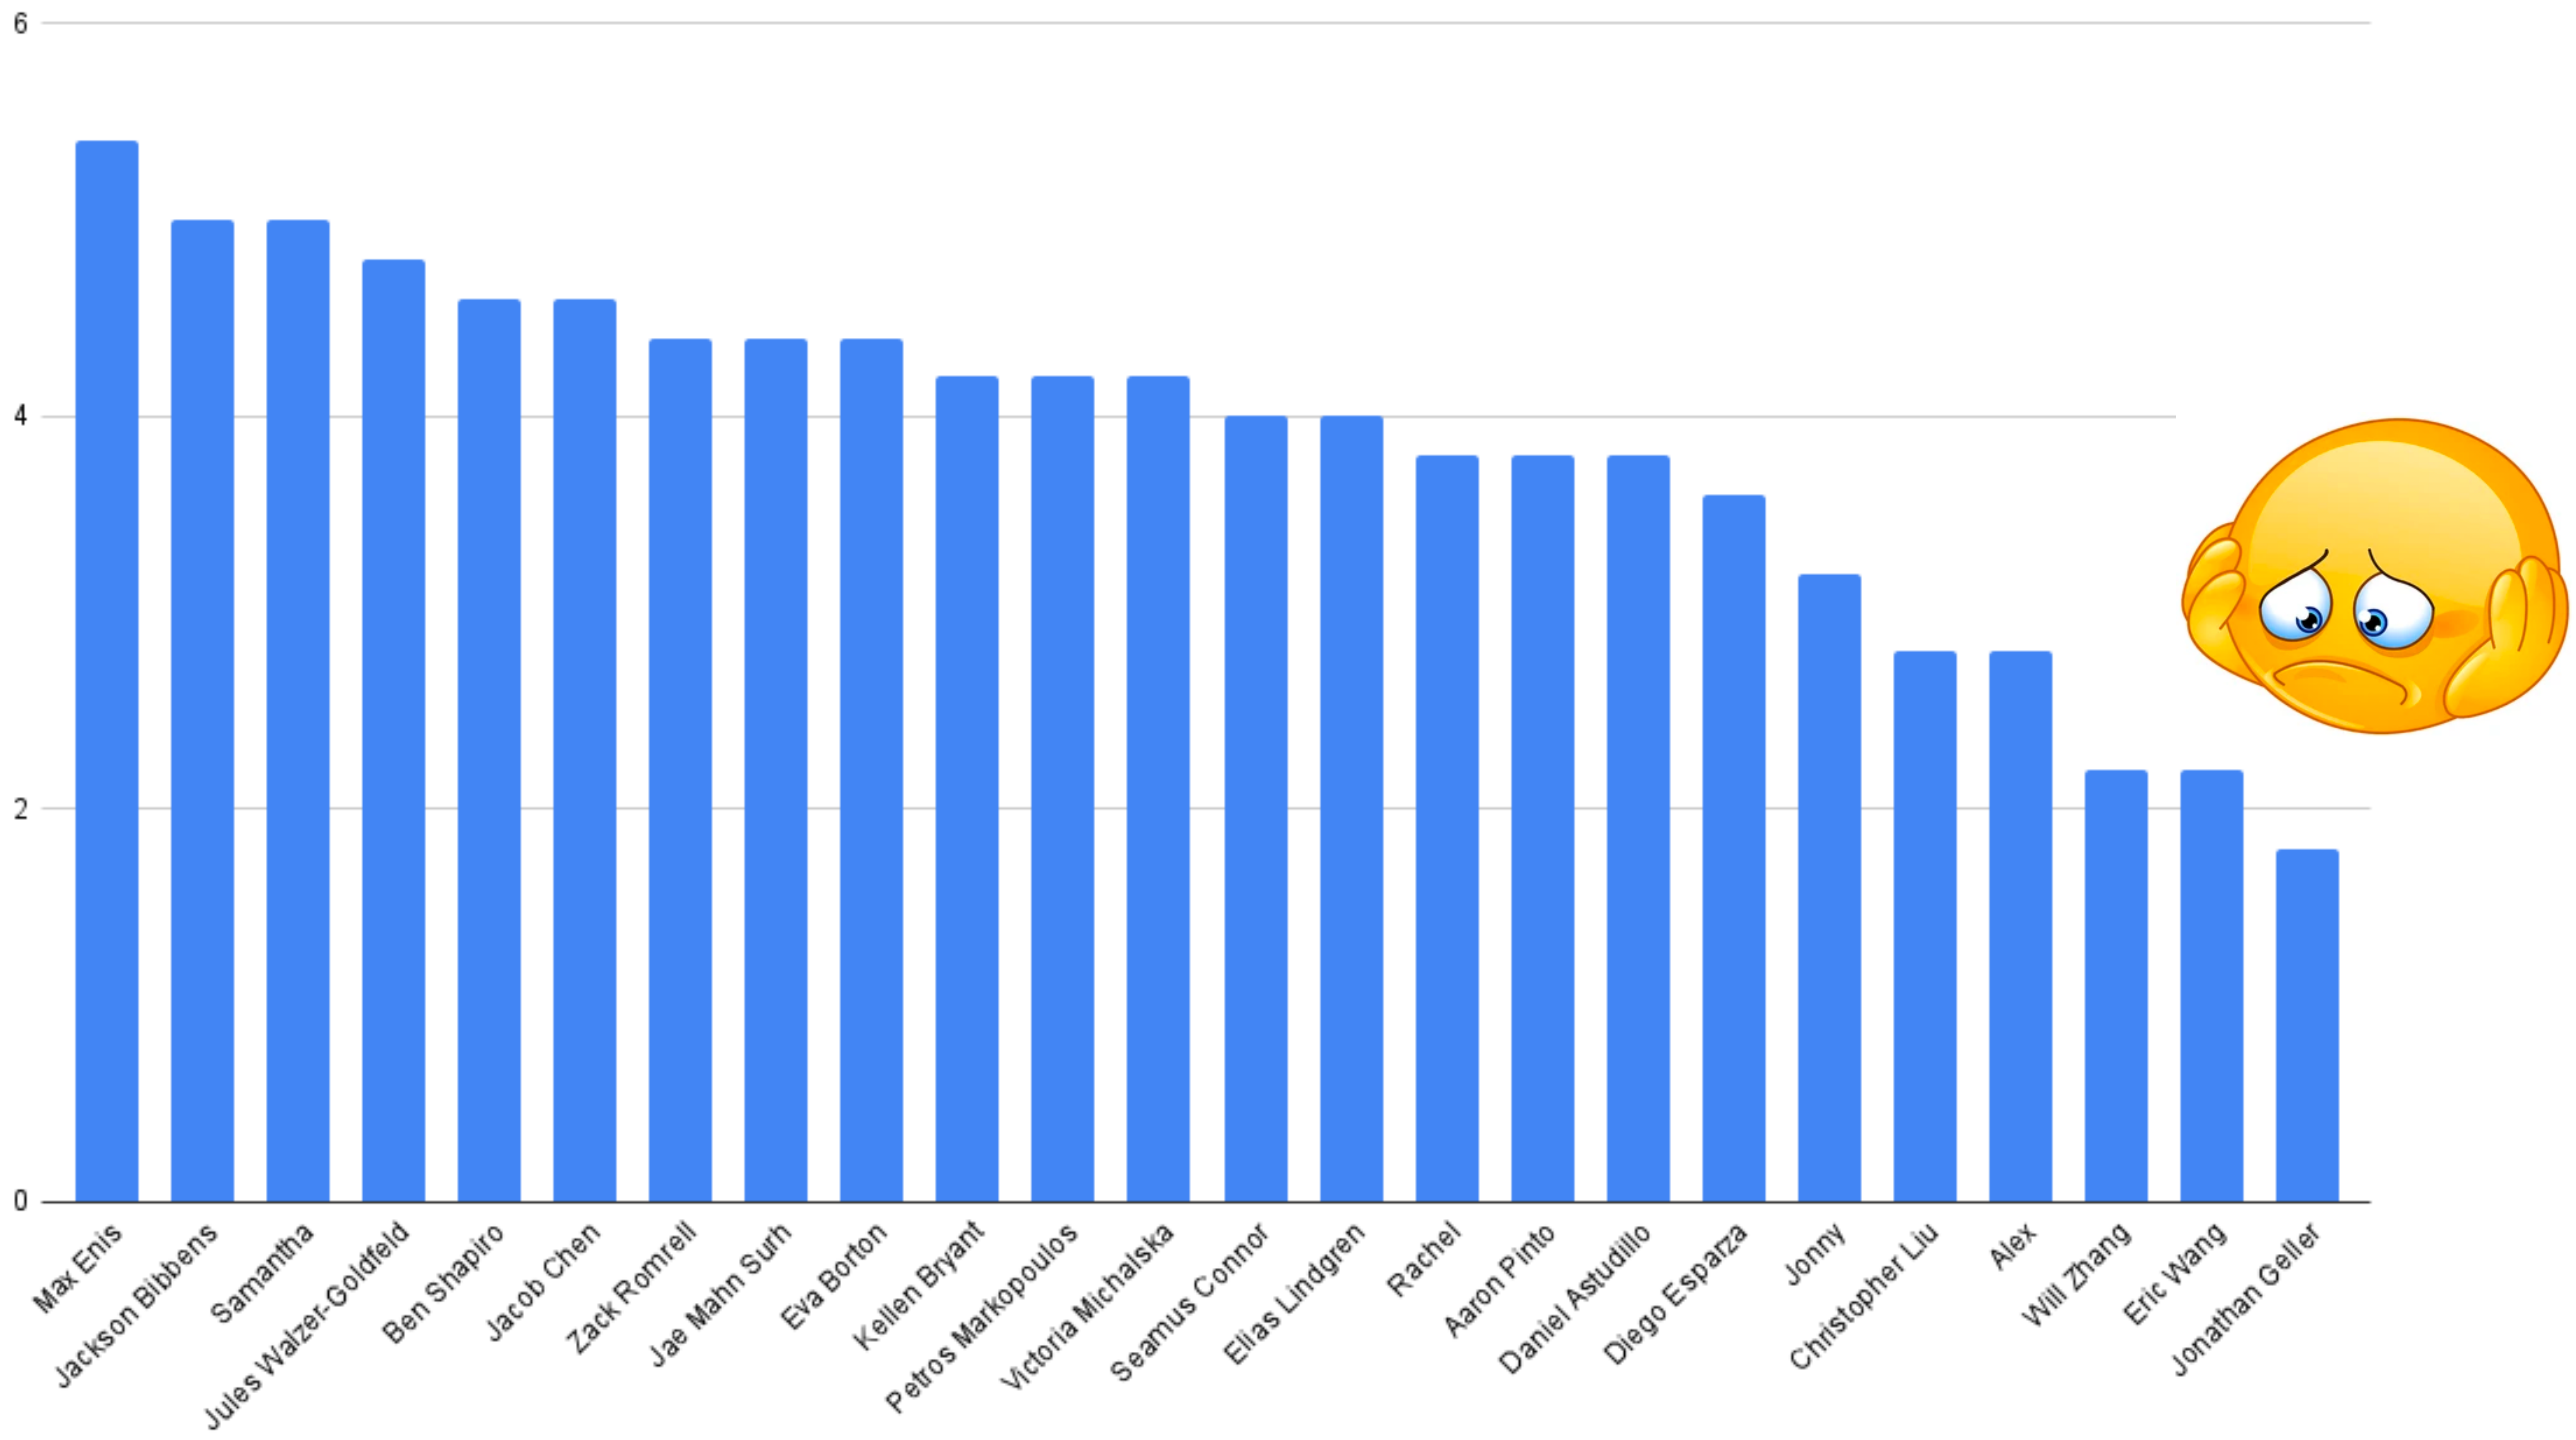

Class First Price Auction

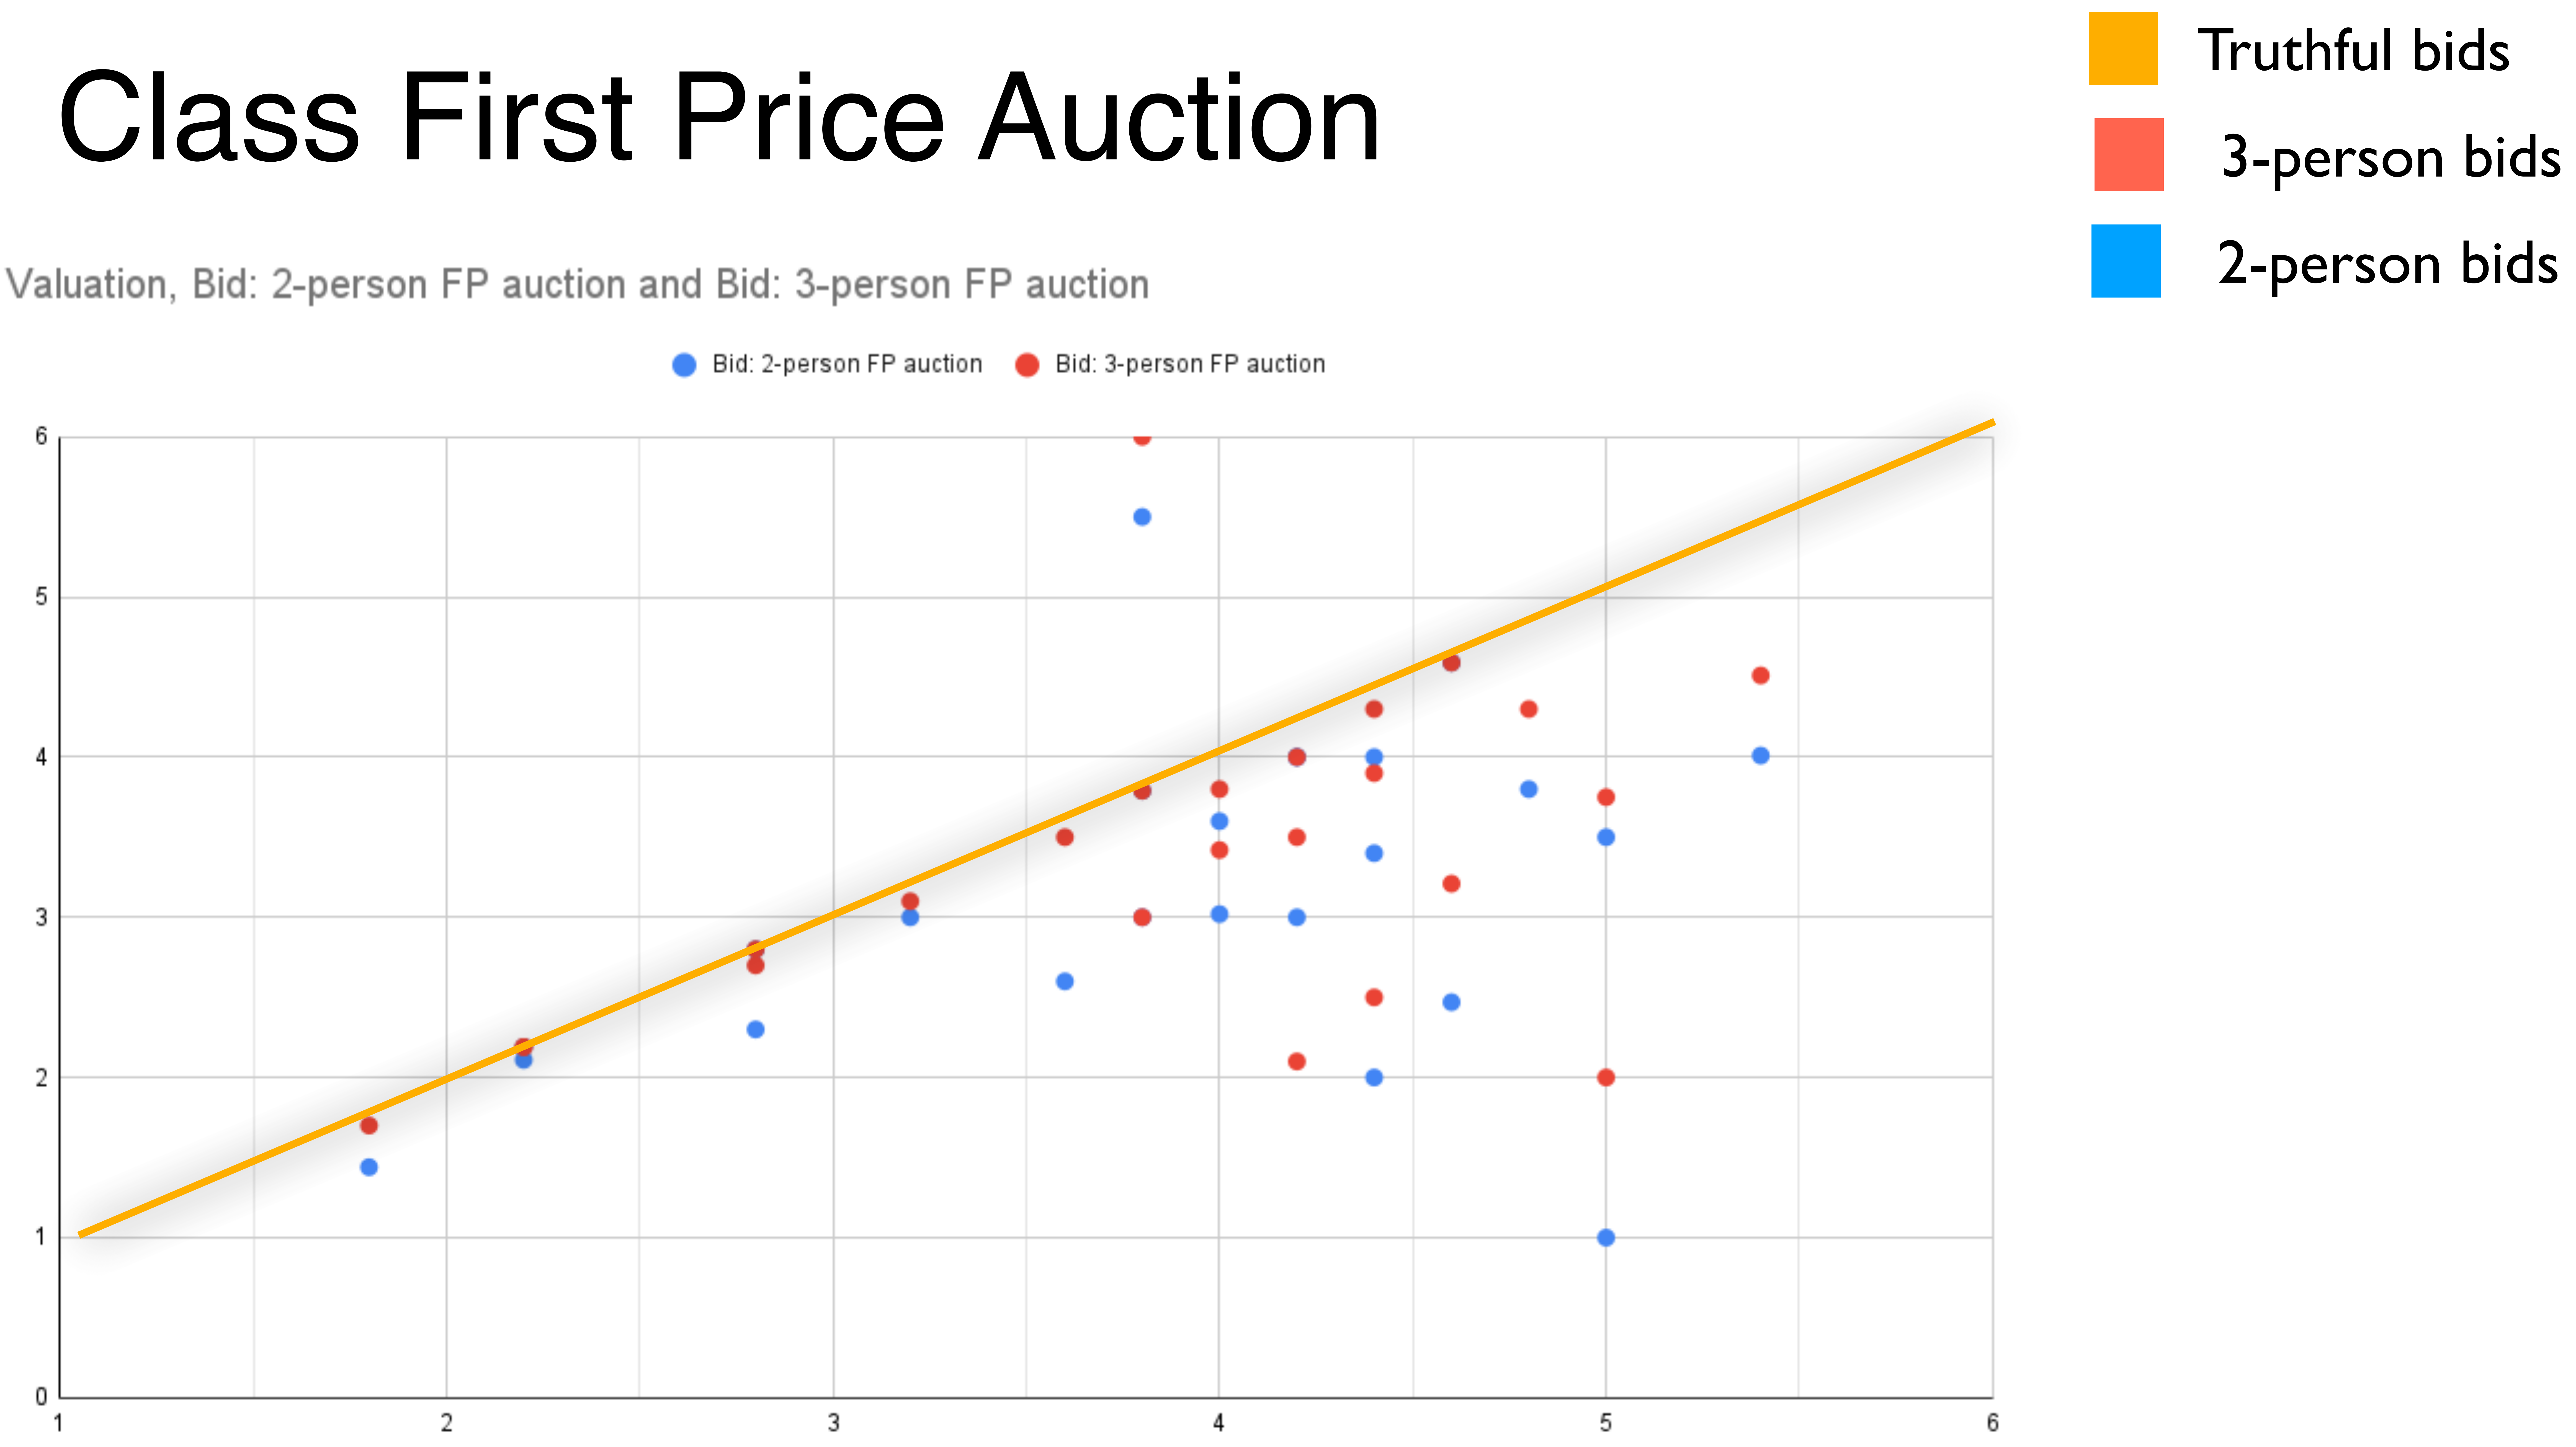

Class First Price Auction

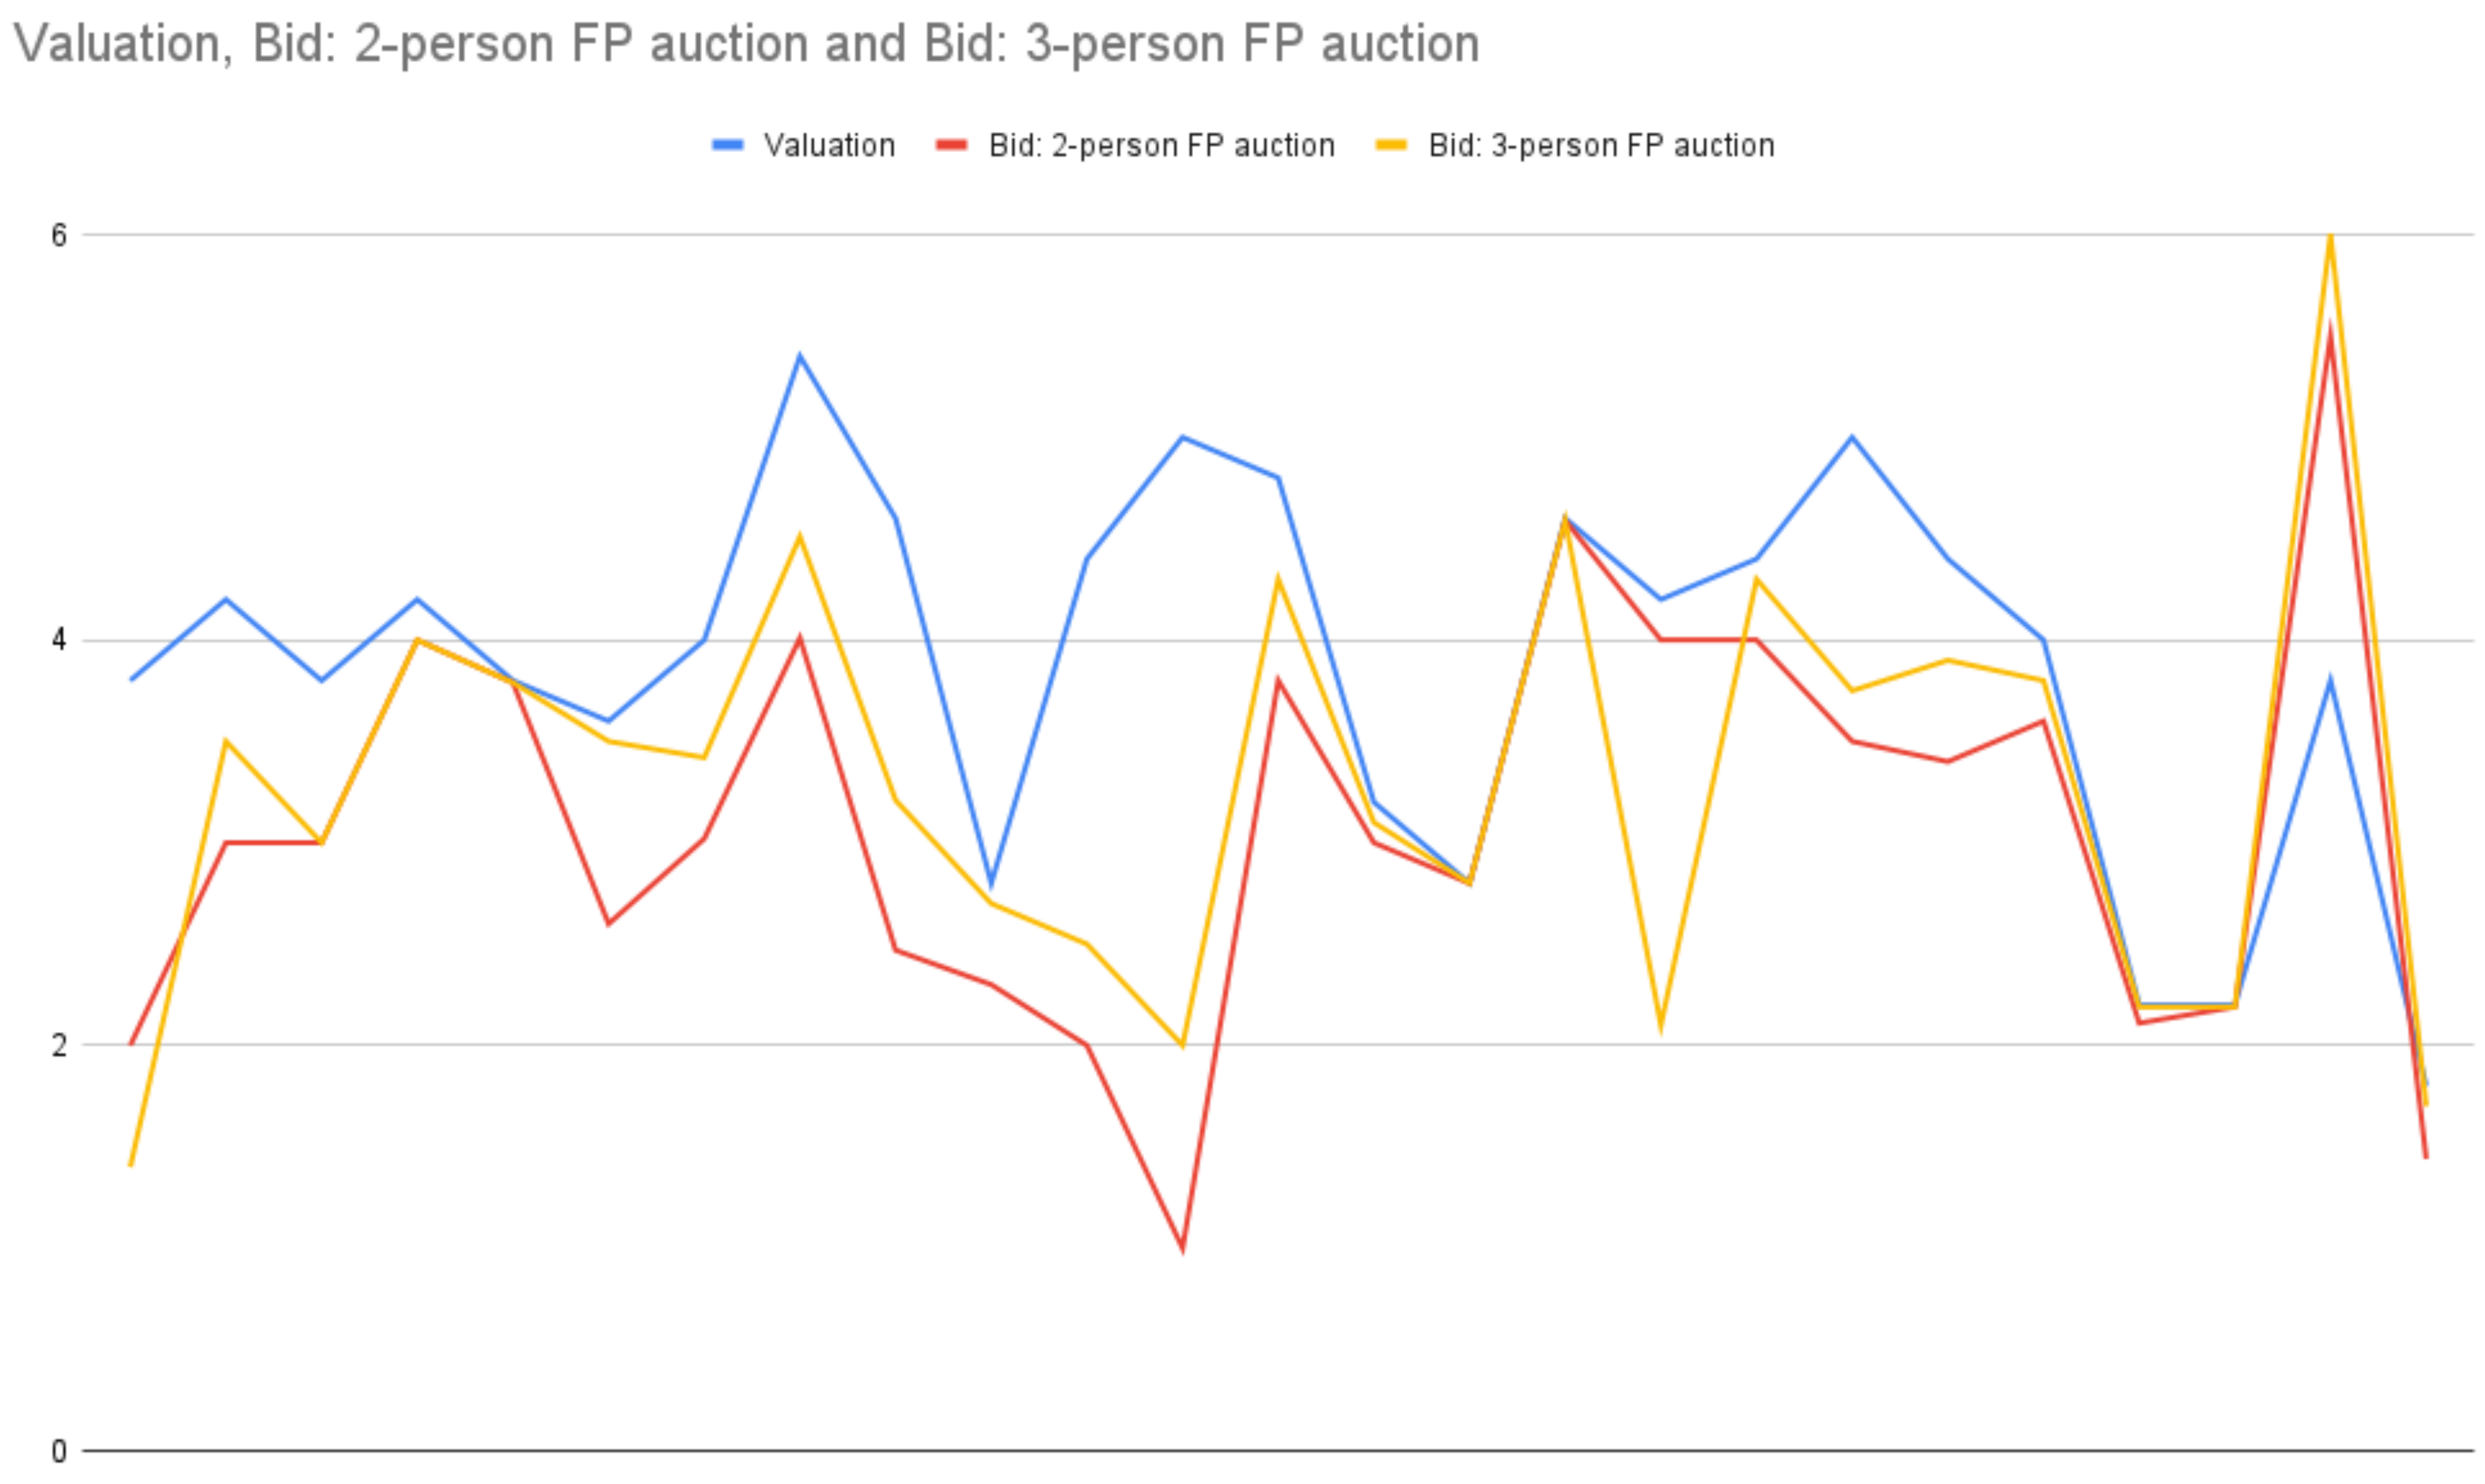

Winners: 2-Bidder Auctions

- Daniel: -1.7$
- Alex, Jae: 0 dollars
- Jacob, Aaron: 0.01$
- Petros: 0.2$
- Victoria: 0.2$
- Seamus: 0.98$
- Eva, Jules: 1$
- Max: 1.39$
- Samantha: 1.5$

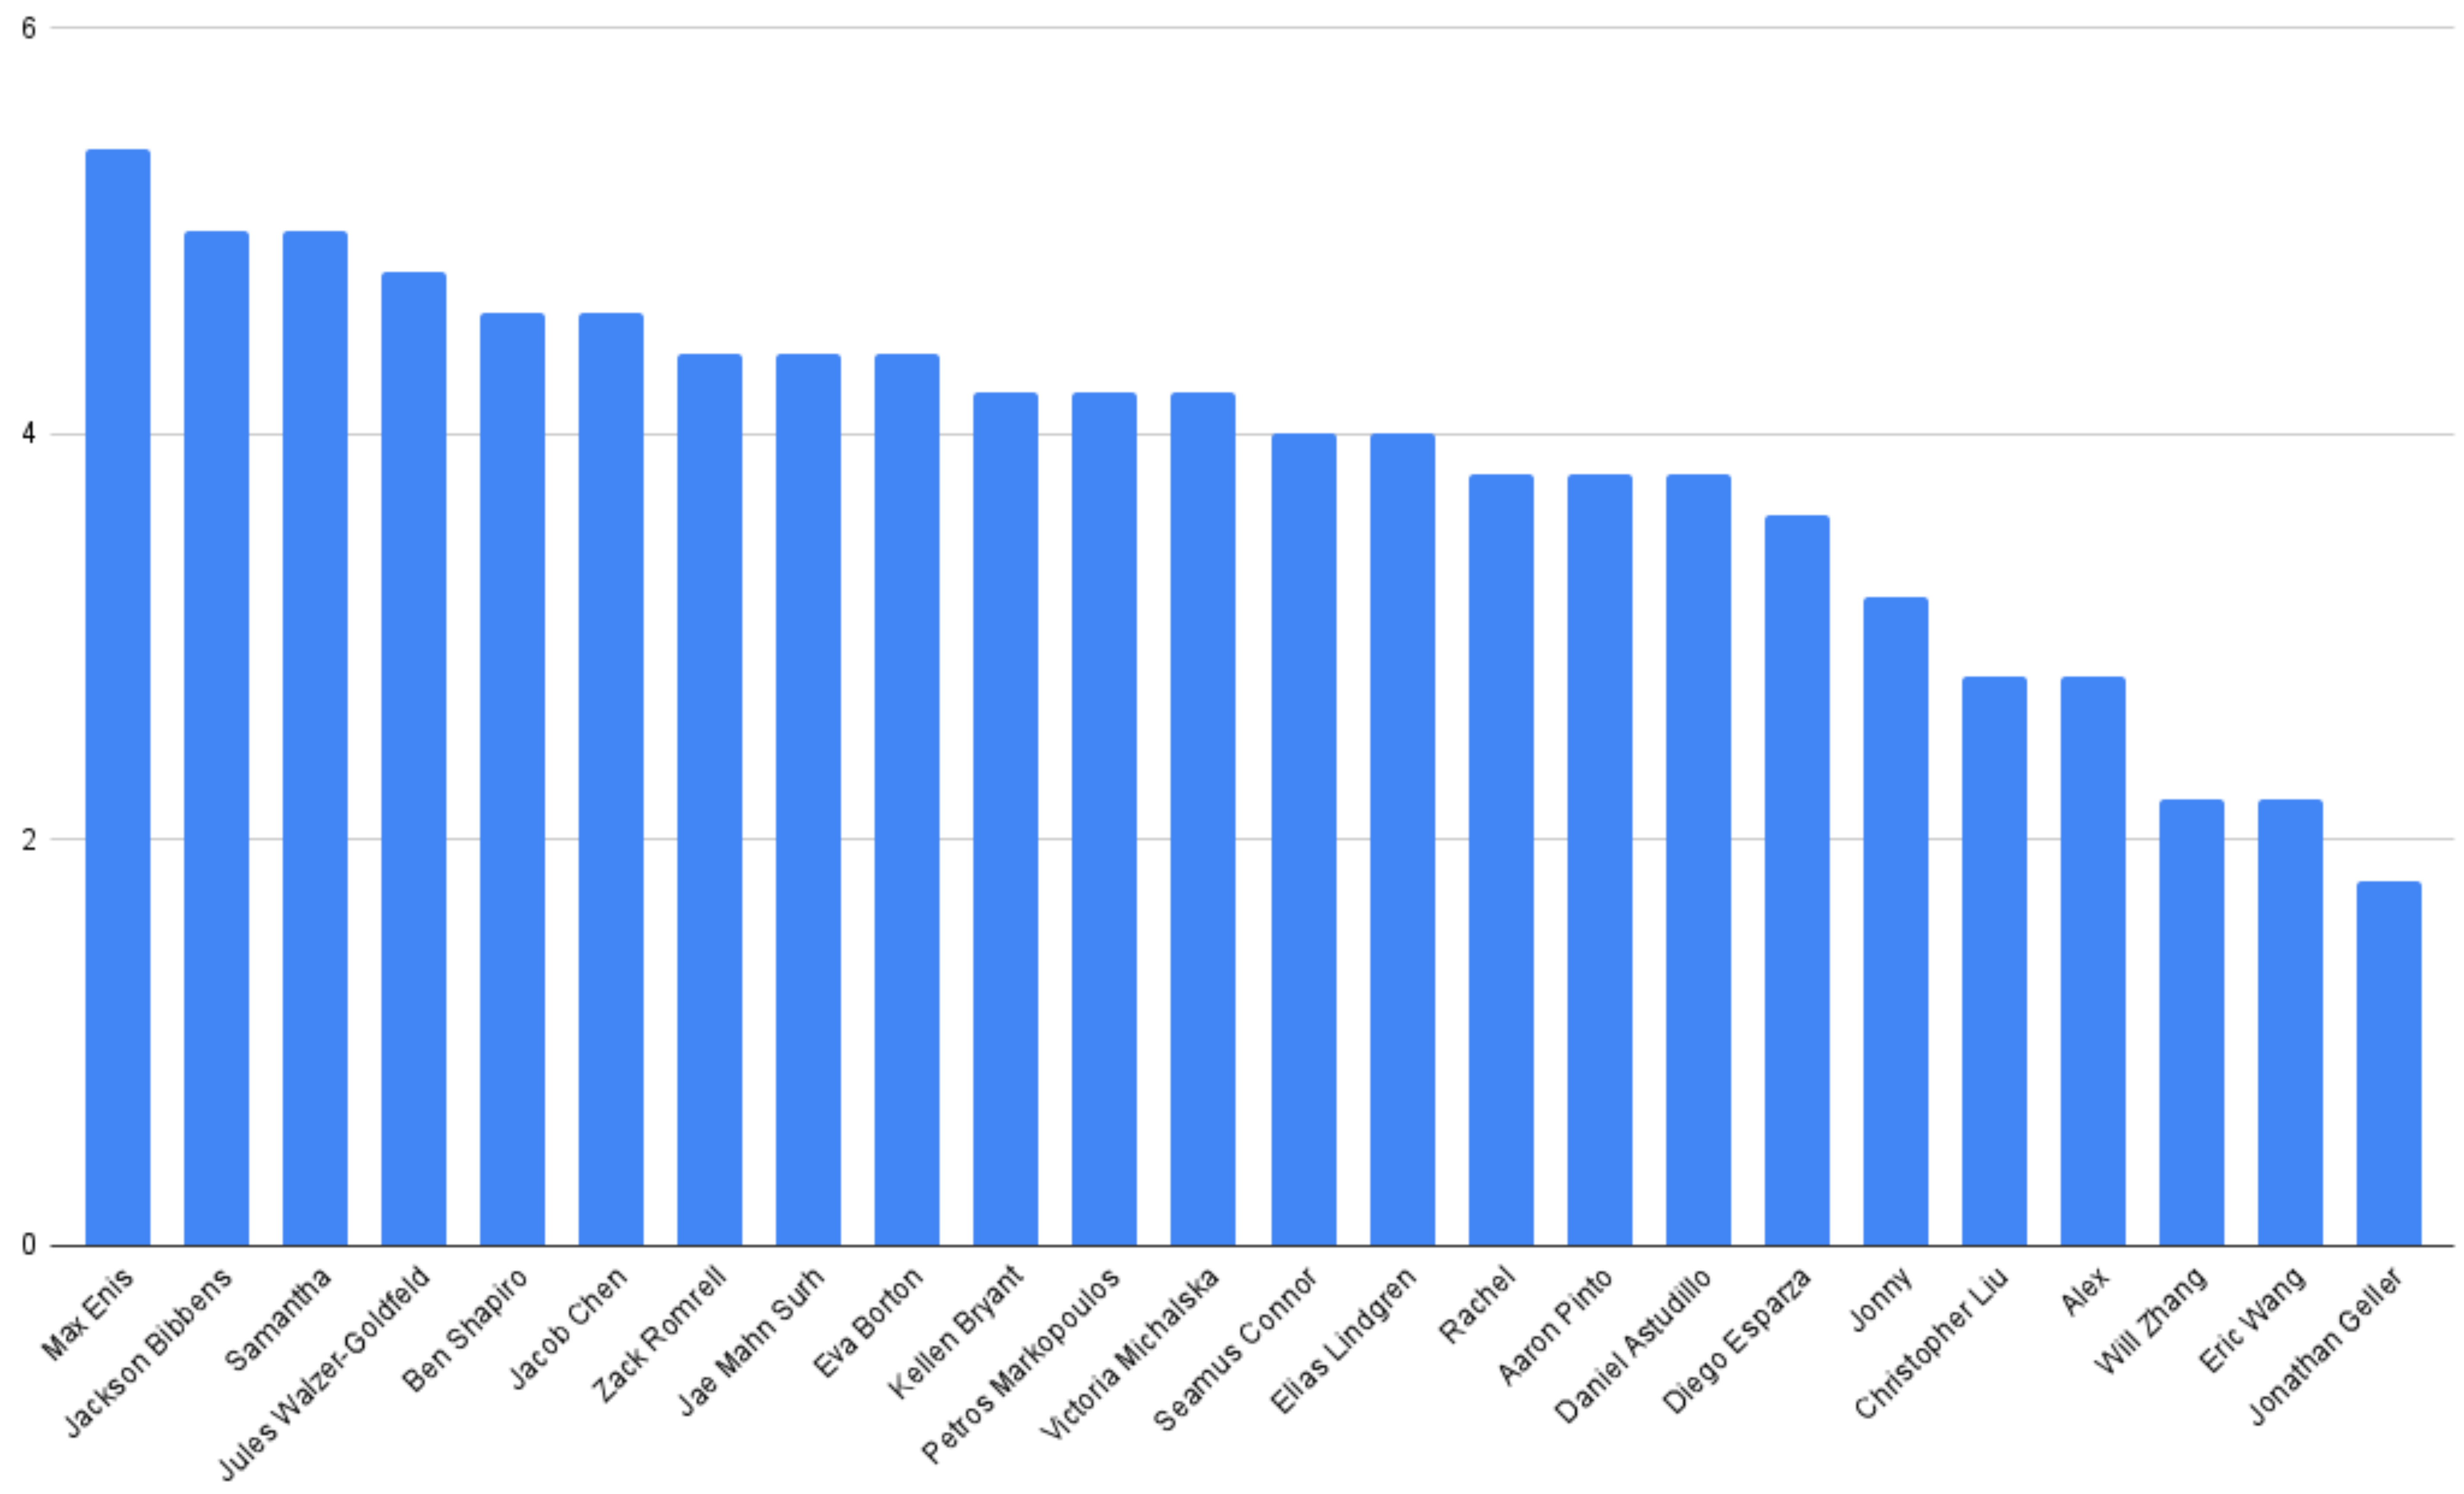

Total winnings: 4.99 $

2 Bidder Auction

Average revenue: 3.9

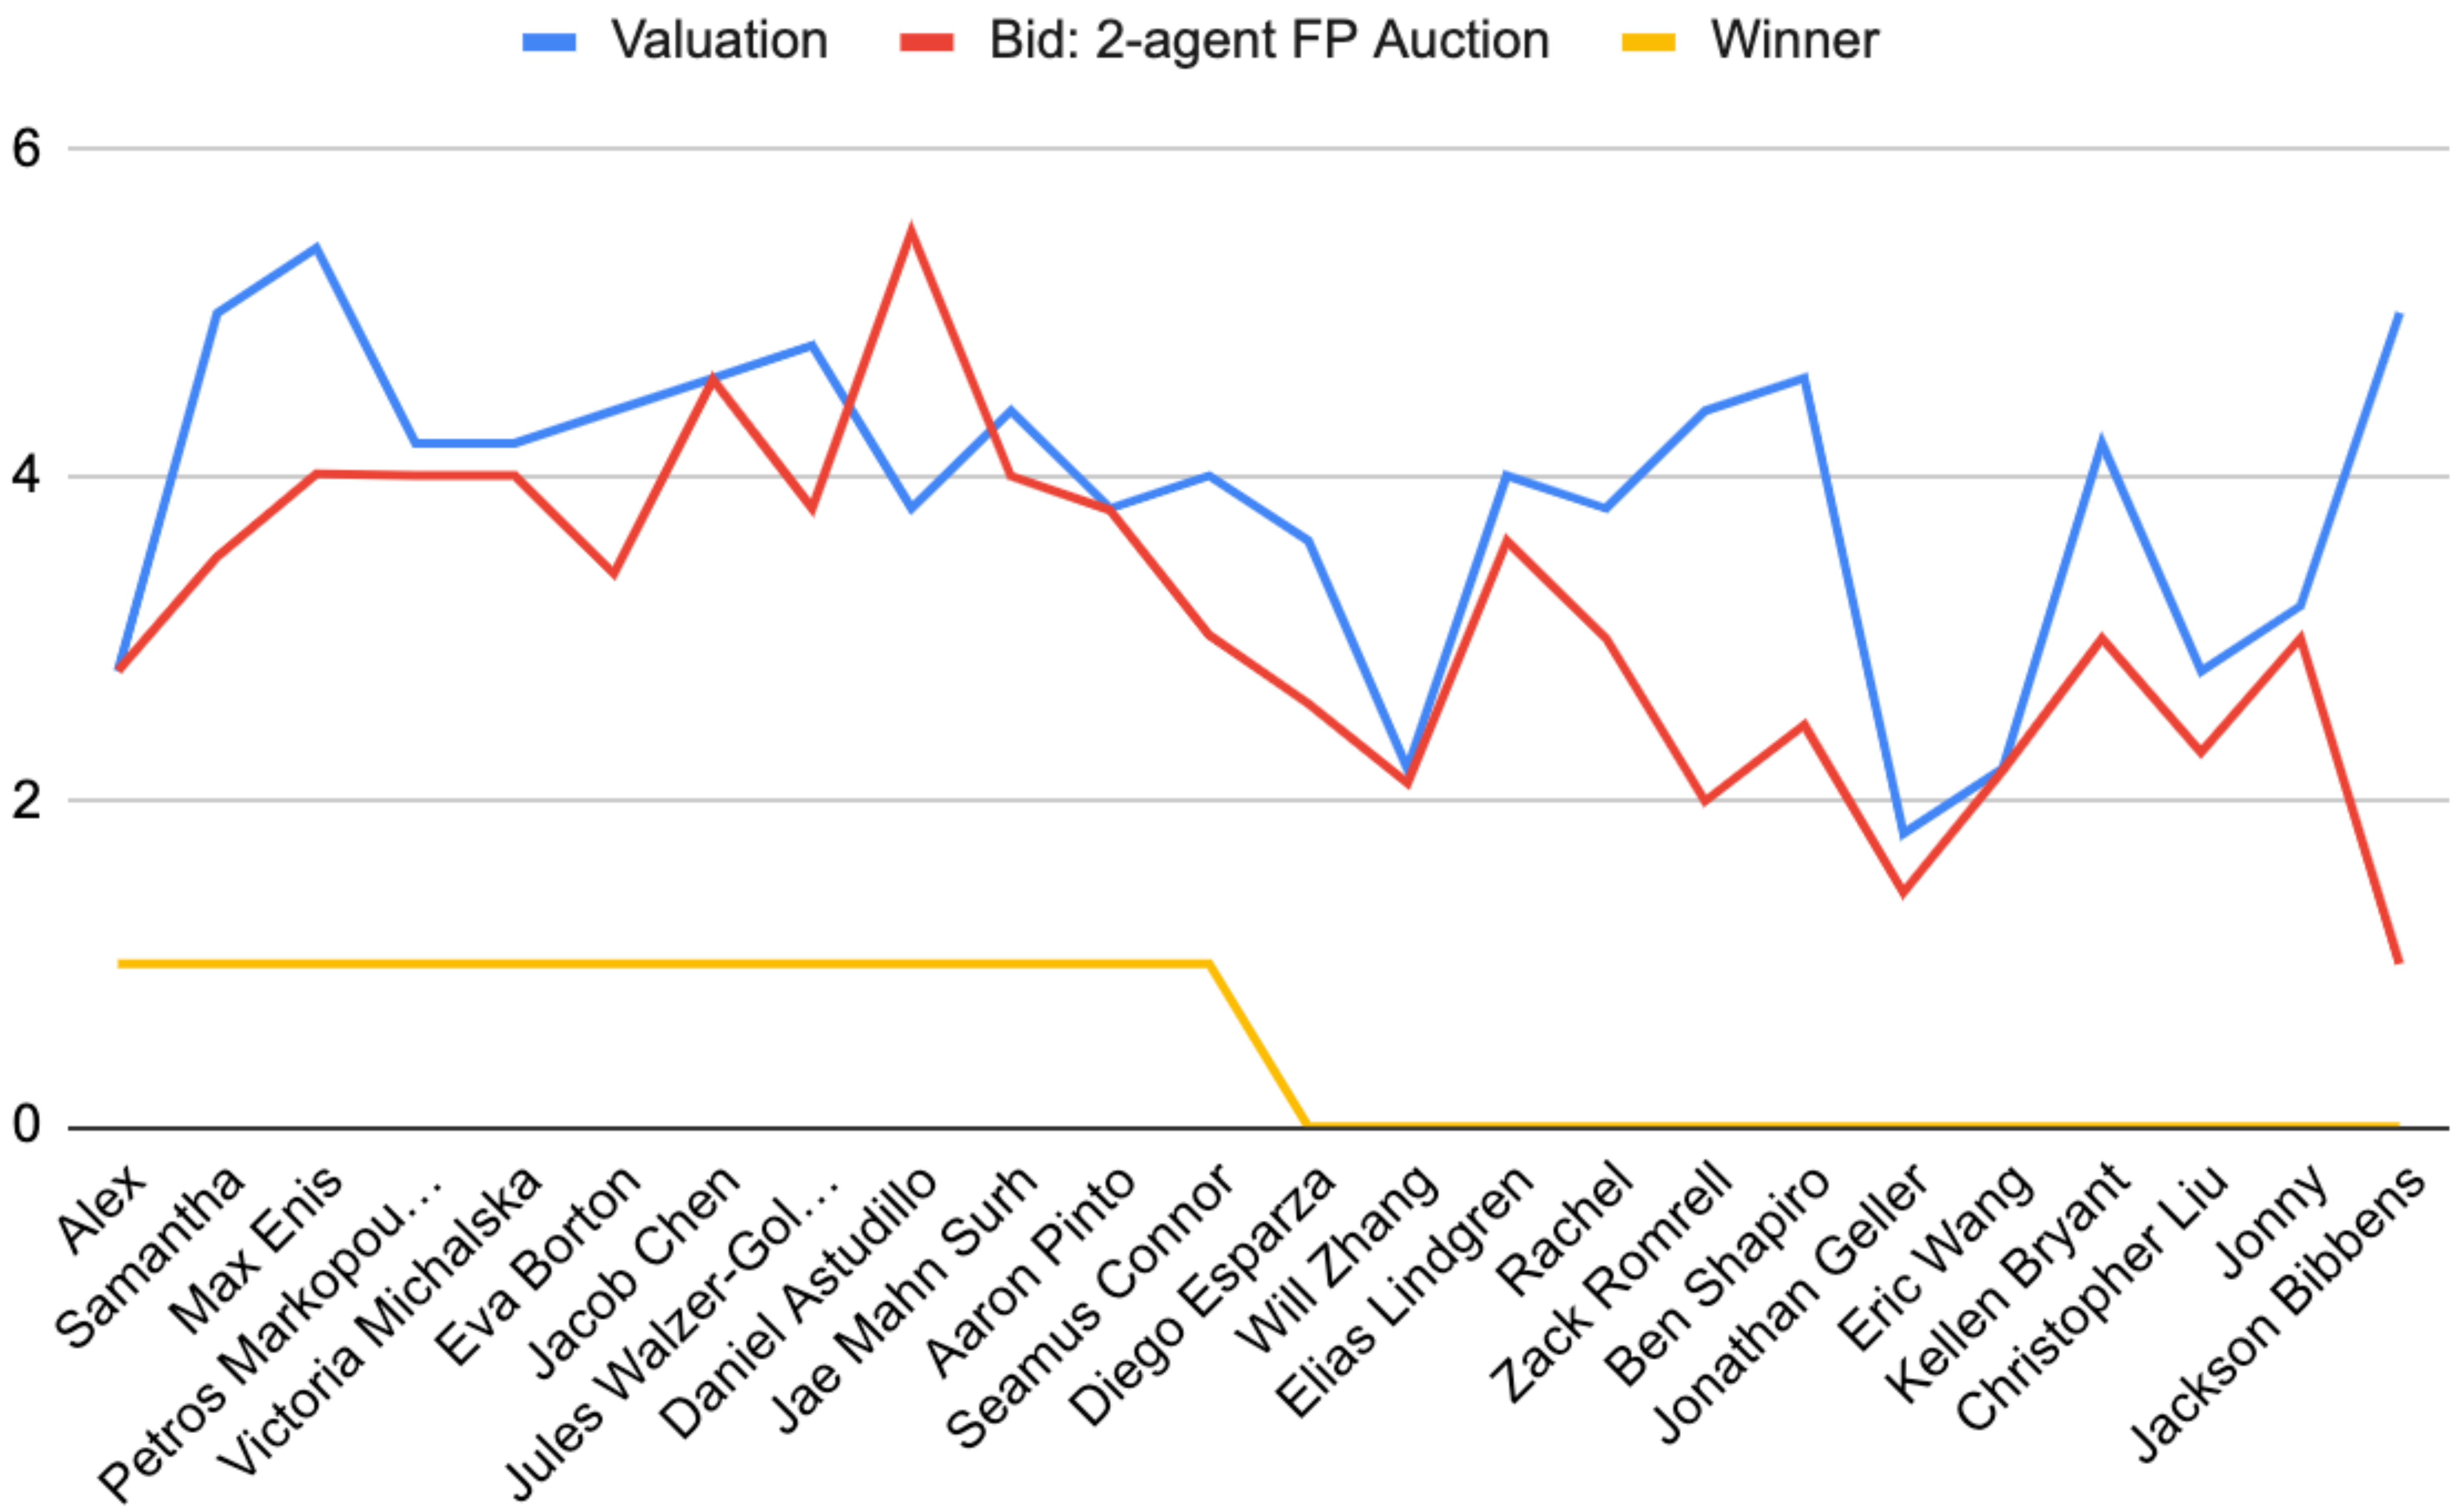

Winners: 3-Bidder Auctions

- Daniel: -2.2$
- Jacob: 0.01$
- Jae: 0.1$
- Petros: 0.2$
- Eva: 0.5$
- Jules: 0.5$
- Max: 0.89$
- Ben: 1.39$

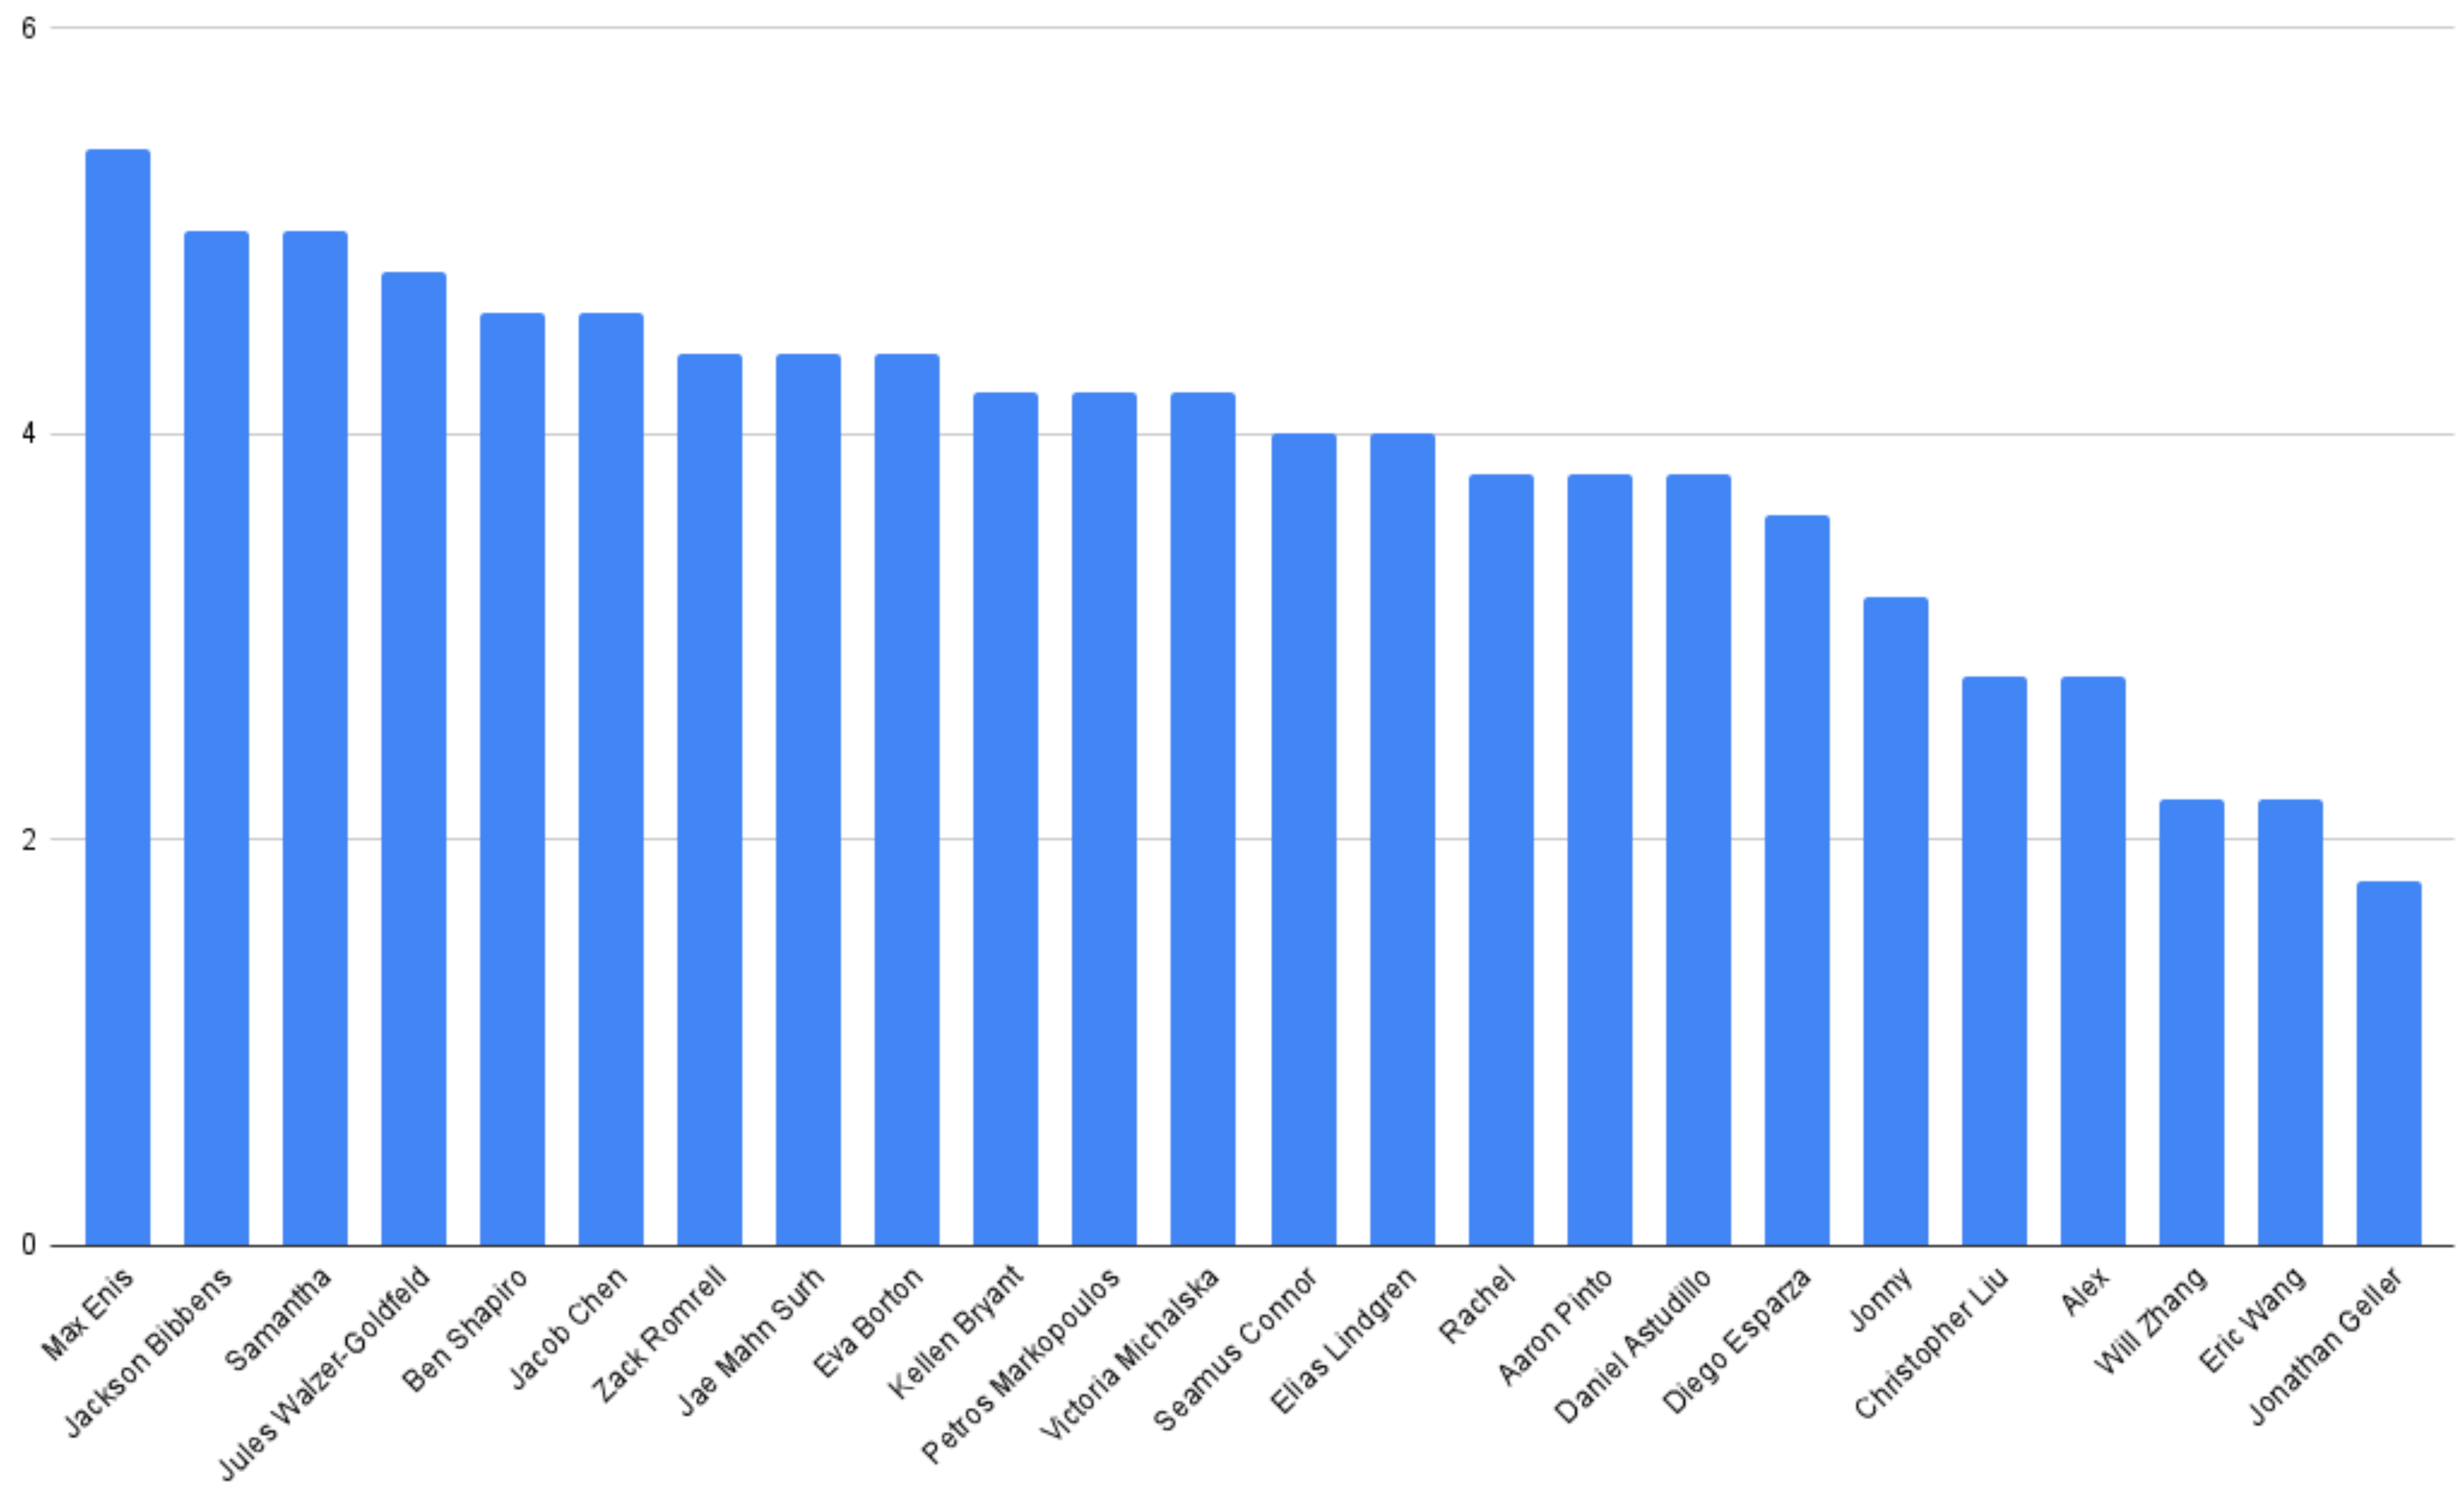

Total winnings: 1.39$

3 Bidder Auction

Average revenue: 4.35

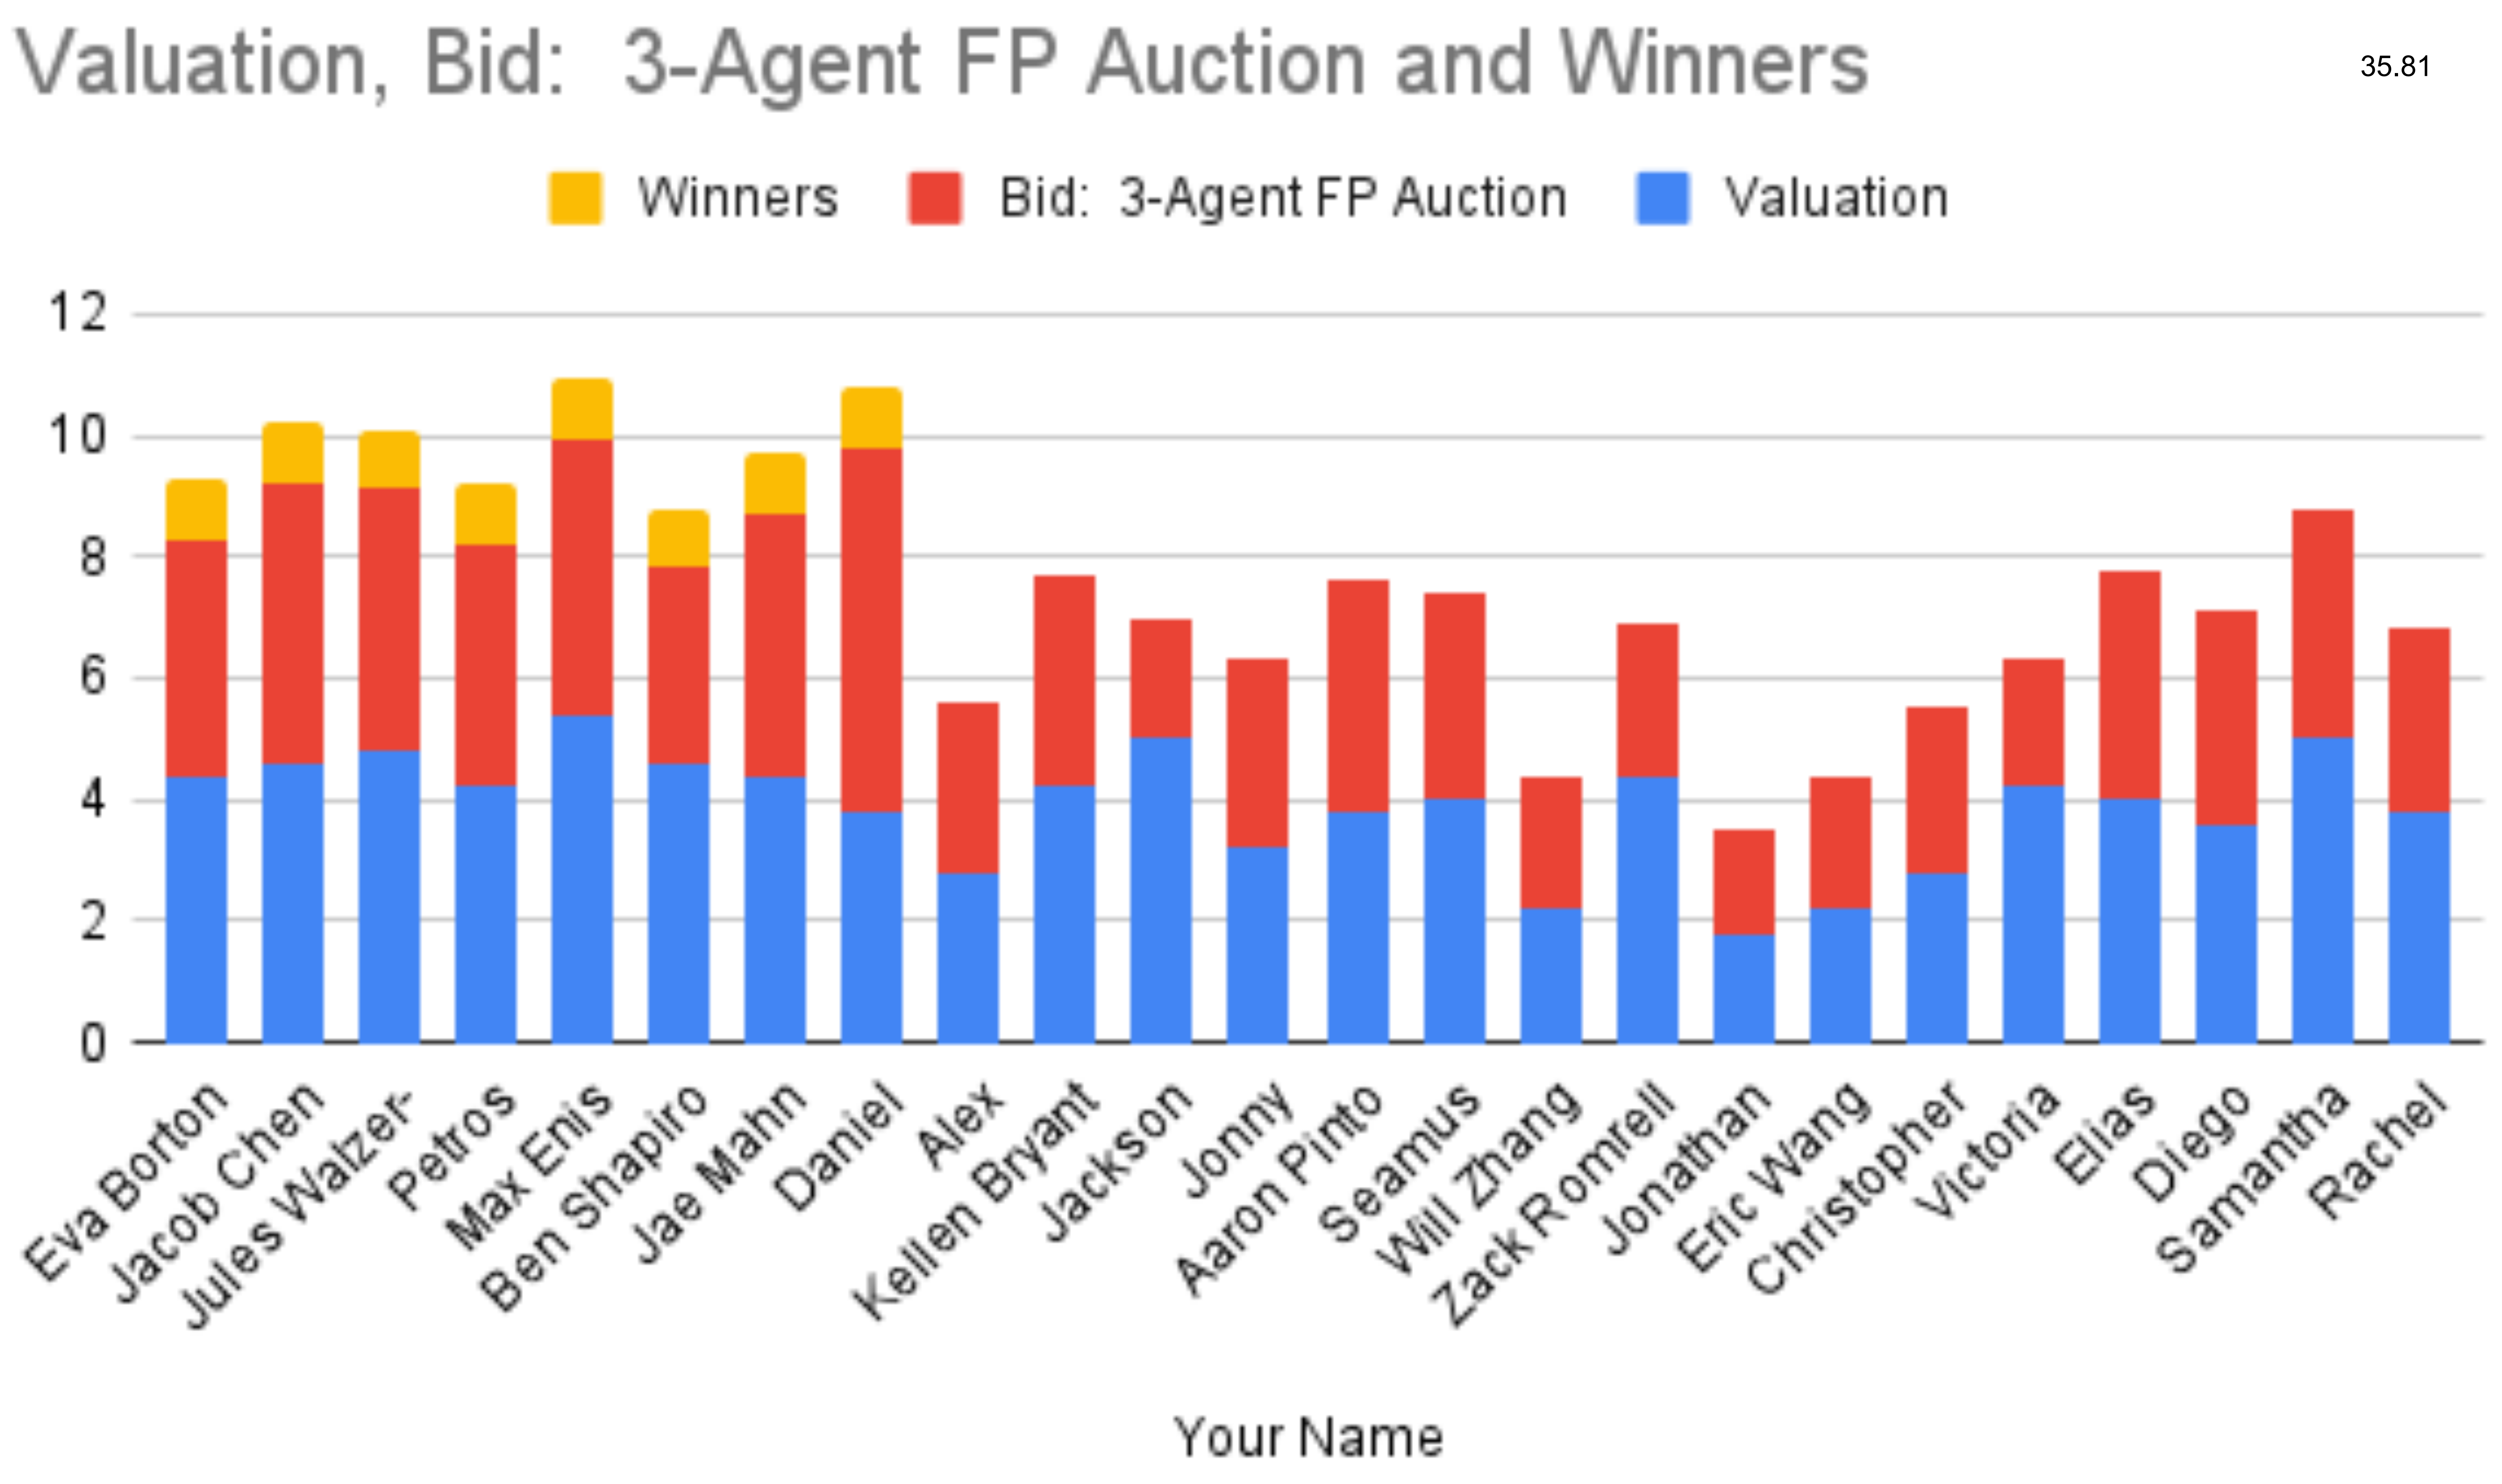

Takeways?

- Bidders will not bid truthfully in first-price auction
  - But rather shade their bid down
- Competition drives the bids up!
  - More number of bidders means more revenue
  - If sellers care about revenue, need to get more participation
- It is difficult for bidders to reason about equilibrium strategies
- **Questions.** What is the theoretical equilibrium that bidders reach?
  - Does our class auction match what theory says?
  - Which auction (first price or second price) generates more revenue?

Credit: Textbook by Parkes and Seuven

First-Price vs Second Price

Both the first-price and second-price auction (at equilibrium) generate the same revenue!

To show this, we need to analyze first-price auction, which is an incomplete-information or "Bayesian game"

Bayesian Nash Equilibrium

Bayesian Games

- In Bayesian games, a player may end up playing different games based on the private types of players
- We make two simplying assumptions:
  - Any private information must pertain only to utilities, that is, all possible games have the same no. of agents and the same strategy space for each agent, differing only in payoffs
  - Players maintain beliefs about the game (i.e., about the utilities) in the form of a probability distribution over types
    - This probability distribution is common knowledge and is called a **common prior**
    - For example, in first-price auctions we will assume that valuation of all agents drawn independently and identically from a distribution (which everyone knows)

After receiving any private information (their own type), players can update their beliefs conditioned on this information (using Bayes rule): hence the name Bayesian games

Bayesian Auction & Assumptions

- Assume bidders have **independent private value (IPV)**
  - Each bidders value $v_i \sim G_i$ is sampled independently, according to a distribution function G_i
  - G_i is continuously differentiable with full support on $[0, v_{\max}]$ so that the probability density function $g_i(z)$ satisfies $g_i(z) > 0$ everywhere
- The distributions $G_1, G_2, \ldots, G_n$ are **common knowledge**
  - Every bidder knows the distributions and knows that others know it as well
- When $G_i = G$ for all bidders and each bidders values are drawn independently and identically from the distribution G
  - We say values are drawn **i.i.d from** G
- For first-price auction: we will assume values are drawn **i.i.d from the uniform distribution** on $[0,1]$

Bayesian Nash Equilibrium

- A strategy or plan of action for each player (as a function of types) should be such that it maximizes each players expected utility
    - expectation is over the types of other players
- Given a Bayesian game with independent private values v_{-i} , i's interim **expected utility** for a strategy profile $s = (s_1, \ldots, s_n)$ is

$$u_i(s \mid v_i) = \sum_{v_{-i} \in G} u_i(s, v_i, v_{-i}) \cdot \text{Pr}(v_{-i})$$

(v_i, v_{-i}) together give full information about utilities to the players

- A strategy profile $s = (s_i, s_{-i})$ is a **pure strategy Bayes Nash equilibrium** if no player can increase their interim expected utility by unilaterally changing their strategy